

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 723 238 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
19.09.2001 Bulletin 2001/38

(51) Int Cl.7: **G06F 17/30**

(21) Application number: **96300429.6**

(22) Date of filing: **23.01.1996**

(54) **Relational database system and method with high data availability during table data restructuring**

Relationales Datenbanksystem und Verfahren mit grosser Verfügbarkeit der Daten bei der Umstrukturierung von Tabellendaten

Système de gestion de base de données relationnelle et procédé avec grande disponibilité de données pendant la restructuration de tables

(84) Designated Contracting States:
DE FR GB IT SE

(30) Priority: **23.01.1995 US 377758**

(43) Date of publication of application:
24.07.1996 Bulletin 1996/30

(73) Proprietor: **Compaq Computer Corporation**
Houston, Texas 77070-2698 (US)

(72) Inventors:
• **Maler, Donald S.**
Campbell, California 95008 (US)
• **Marton, Roberta S.**
Fremont, California 94539 (US)
• **Troisi, James H.**
Sunnyvale, California 94087 (US)

• **Celis, Pedro, (NMI)**
Austin, Texas 78759-6123 (US)

(74) Representative:
Cross, Rupert Edward Blount et al
BOULT WADE TENNANT,
Verulam Gardens
70 Gray's Inn Road
London WC1X 8BT (GB)

(56) References cited:
EP-A- 0 306 197 **EP-A- 0 560 543**

• **INTELLECTUAL LEVERAGE, SAN FRANCISCO,**
FEB. 25 - MAR. 1, 1991, no. CONF. 36, 25
February 1991, INSTITUTE OF ELECTRICAL
AND ELECTRONICS ENGINEERS, pages
105-109, XP000293859 LESLIE H: "OPTIMIZING
PARALLEL QUERY PLANS AND EXECUTION"

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

[0001] The present invention relates generally to database management systems, and particularly to a database management system in which database table availability is maintained, with minimal or no user availability outages during table restructuring operations such as splitting a table or index partition, moving an existing table or index partition, creating a new index and moving a table or index partition boundary.

BACKGROUND OF THE INVENTION

[0002] Database configuration and reconfiguration operations can have a significant effect on the availability of user applications that need access to databases undergoing structural changes. The Tandem™ NonStop™ SQUMP relational database management system (DBMS), prior to the present invention, allowed read access, but not write access, to the portions of the database table undergoing the restructuring operation.

[0003] Although most users perform these operations infrequently, their duration can account for thousands of minutes of application outages per year. A discussion of the cost of application outages appears in the article "An Overview of NonStop SQUMP," Ho et al., Tandem Systems Review, July 1994.

[0004] Embodiments of the present invention eliminate most of the downtime associated with four database reconfiguration operations: Move Partition, Split Partition, Move Partition Boundary, and Create Index. The Move Partition procedure moves a partition that resides on one disk to another disk. The Split Partition procedure splits one partition into two.

The Move Partition Boundary procedure moves rows of a base table between adjacent partitions, typically shifting rows from a large partition to a less large partition. The Create Index procedure creates an efficient alternate access path to a database table by ordering data according to the value specified in the key columns. The Create Unique Index variant of this procedure ensures that there is exactly one alternate access path to each record in the database table.

[0005] The implementation of these procedures in embodiments of the invention substantially reduces, but does not eliminate, associated outages. Even with the present invention, user database activity continues to be restricted for about one minute or less per database restructuring operation. The outage time varies depending on the number of user transactions running against the table being restructured, the size of those transactions, and the number of partitions in the affected table.

[0006] Embodiments of the present invention enable substantially improved user transaction access to a database table while the table is undergoing a structural change operation.

[0007] Embodiments of the present invention also enable Move Partition, Split Partition, Move Partition Boundary, and Create Index operations to be performed on a database table while enabling user transactions to continue to be performed on the table, except during a short final phase of those operations.

[0008] Embodiments of the present invention also avoid the use of "side files" for the above mentioned database reconfiguration operations, so as to reduce the overhead disk space requirements associated with these operations to almost zero.

[0009] Embodiments of the present invention also implement the above mentioned database reconfiguration operations using procedures that directly read from the transaction audit log, thereby making these database reconfiguration operations similar to a database recovery process.

[0010] Embodiments of the present invention also implement the database reconfiguration operations using well established, optimized, database recovery process procedures thereby making the database reconfiguration operations efficient in terms of computation time and in terms of the computer resources used. Embodiments of the present invention also utilize substantially similar database reconfiguration procedures for four distinct database reconfiguration operations, thereby improving the reliability of those procedures and simplifying maintenance of and updates to those procedures.

SUMMARY OF THE INVENTION

[0011] Accordingly, the present invention provides a computer system for storing and providing user access to data in stored database objects, comprising: memory for storing said database objects, said database objects comprising database tables, database partitions and database indexes, wherein said memory resides in a plurality of interconnected computer nodes; a central processing unit; a transaction manager, coupled to said memory, for managing computational transactions that add, delete and alter data stored in said database objects; said transaction manager including audit trail generation instructions for generating an audit trail and storing said audit trail in said memory, said audit trail including audit records, at least a subset of said audit records each denoting an event selected from the set consisting essentially of addition, deletion and alteration of specified data in a specified one of said database objects; a restructuring procedure, executable by said central processing unit, for restructuring a specified one of said database

objects;

characterized in that said restructuring procedure includes:

first phase instructions for accessing data in a first of said objects and storing corresponding data in a second of said objects while allowing continued performance of computational transactions against said first object; second phase instructions, for execution after said first phase instructions, for accessing said audit records in said audit trail created by said transaction manager during execution of said first phase instructions and thereafter until execution of said second phase instructions is completed; said second phase instructions updating said data stored in said second object by redoing with respect to said second object each event denoted by said accessed audit trail records; and third phase instructions, for execution after said second phase instructions, for obtaining a lock on said first object so as to prevent continued performance of computational transactions against said first object, and for then accessing audit records in said audit trail created by said transaction manager after execution of said second phase instructions; said third phase instructions updating said data in said second object by redoing with respect to said second object each event denoted by said accessed audit trail records.

[0012] A preferred embodiment of the invention comprises a set of procedures for modifying the structure of a database table or index "online," while the database table or index remains available for execution of transactions, with minimal impact on the availability of the database table for transaction execution.

[0013] The preferred embodiment operates in a database computer system having memory, residing in a plurality of interconnected computer nodes, for storing database tables. Each database table or index has a plurality of columns, a primary key index based on a specified subset of the columns, and an associated table schema. In most implementations, at least some of the database tables or indexes are partitioned into a plurality of partitions, each partition storing records having primary key values in a primary key range distinct from the other partitions.

[0014] A transaction manager generates and stores an audit trail, each audit entry denoting a database table or index record event, such as an addition, deletion or alteration of a specified database table or index record in a specified one of the database tables or indexes.

[0015] Four online data definition procedures allow the structure of a database table or index to be altered while the database table or index remains available for execution of transactions, with minimal impact of the availability of the database table or index for transaction execution. The four online data definition procedures are a Move Partition procedure, a Split Partition procedure, a Move Partition Boundary procedure, and a Create Index procedure. Each of these online procedures has several phases of execution. In a first phase, the definitions of existing objects are read and any new objects that are needed to perform the requested procedure are created. Then, records of a table or index partition or the entire table are accessed, using "browse" access, so as to generate a new partition, to move records between two partitions, or to create a new index. Browse access is a form of read access that allows database records to be read through any record locks that may be in place. Thus, browse access will sometimes result in the reading of records that are in the midst of being modified.

[0016] In a second phase, audit trail entries are accessed and the equivalent database table or index operations denoted in those audit trail entries are redone on the target objects, whenever necessary, to bring the data records created during the first phase up-to-date. In a third phase, access to the database table is briefly locked (A) while audit trail entries created after the second phase are used to make final changes to the previously created data records, and then (B) while the database table or index schema is updated to reflect the changes to the database table or index produced.

[0017] In a fourth phase, used by the Move Partition Boundary and Split Partition procedures, records in a database partition that are inconsistent with the modified database table schema are deleted as a background operation while use of the database table by transactions resumes.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings, of which:

Figure 1 is a block diagram of a computer system having a database management system in accordance with the present invention.

Figure 2A is a block diagram of a database table. Figure 2B is a block diagram of an alternate index. Figure 2C is a block diagram of the data structure of an Audit Trail.

Figure 3 is a conceptual diagram of the three primary phases of the database restructuring procedures in the present invention.

Figure 4 is a flowchart of the Move Partition procedure used in a preferred embodiment of the present invention.

Figure 5 is a flowchart of the Split Partition procedure used in a preferred embodiment of the present invention.

Figure 6 is a flowchart of the Move Partition Boundary procedure used in a preferred embodiment of the present invention.

Figure 7 is a flowchart of the Create Index procedure used in a preferred embodiment of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0019] Referring to Figure 1, there is shown a computer system 100 for storing and providing user access to data in stored databases. The system 100 is a distributed computer system having multiple computers 102, 104, 106 interconnected by local area and wide area network communication media 108. The system 100 generally includes at least one database server 102 and many user workstation computers or terminals 104, 106.

[0020] When very large databases are stored in a system, the database tables will be partitioned, and different partitions of the database tables will often be stored in different database servers. However, from the viewpoint of user workstation computers 104, 106, the database server 102 appears to be a single entity. The partitioning of databases and the use of multiple database servers is well known to those skilled in the art.

[0021] As shown in Figure 1, the database server 102 includes a central processing unit (CPU) 110, primary memory 112, a communications interface 114 for communicating with user workstations 104, 106 as well as other system resources not relevant here. Secondary memory 116-1, 116-2, typically magnetic disc storage, in the database server 102 stores database tables 120, database indices 122, a database management system (DBMS) 124 for management of the database tables and associated data structures and resources, and one or more catalogs 126 for storing schema information about the database tables 120 as well directory information for programs used to access the database tables. The DBMS 124 includes an SQL executor 128 for executing SQL statements (i.e., database queries) and an SQL catalog manager 130 for maintenance of the catalogs 126 and for performing database definition and restructuring operations. The SQL catalog manager includes the "online DDL procedures" 132 of the present invention for restructuring database tables while providing improved user transaction access to the affected tables.

[0022] The database server 102 further includes a transaction manager 134 for managing transactions, and application programs 136 that are utilized by users to perform transactions that utilize the database tables 120. The transaction manager creates audit entries for each transaction, which are durably stored in an audit trail file 138 in secondary memory.

[0023] End user workstations 104, 106, typically include a central processing unit (CPU) 140, primary memory 142, a communications interface 144 for communicating with the database server 102 and other system resources, secondary memory 146, and a user interface 148. The user interface 148 typically includes a keyboard and display device, and may include additional resources such as a pointing device and printer. Secondary memory 146 is used for storing computer programs, such as communications software used to access the database server 102. Some end user workstations 106 may be "dumb" terminals that do not include any secondary memory 146, and thus execute only software downloaded into primary memory 142 from a server computer, such as the database server 102 or a file server (not shown).

Glossary

[0024] To assist the reader, the following glossary of terms used in this document is provided.

[0025] **SQL:** SQL stands for "Structured Query Language." Most commercial database servers utilize SQL. Any program for accessing data in a database that utilizes SQL is herein called an "SQL Program." Each statement in an SQL program used to access data in a database is called an "SQL statement."

[0026] **Object(s):** An object is a file, database table or other encapsulated computer resource accessed by a program as a unitary structure. In the context of the preferred embodiment, objects are database tables or indexes. In other implementations of the present invention, objects may be other encapsulated computer resources that the end user accesses indirectly through validated methods (i.e., programs) designed specifically to access those computer resources.

[0027] **DDL Statement:** a data definition language statement. DDL statements are used to create and modify database tables.

[0028] **End user:** a person using a workstation to access database information in a database server. End users typically do not have the authority to modify the structure of database tables.

[0029] **Operator:** a person using a workstation who has the authority and access rights to modify the structure of

database tables and to manually initiate compilation of SQL source code programs.

Audit Trail, Database Table and Index Data Structures

[0030] Figure 2A shows the data structure of a typical database table 120-1. The table 120-1 includes a file label 160, which is essentially a compactly stored copy of the catalog information for the database table, which represents the table's schema as well as other information not relevant here. Next, the table includes a primary key B-tree 162. The use of B-trees in database files is well known to those of ordinary skill in the art of database management systems. Next, the table has a data array 163 organized in rows and columns. The rows of the table are often called "records" 164.

[0031] In the context of preferred embodiments of the present invention, every database table has a primary index. The value of the primary index for a particular record is called the primary key, and the primary key is typically equal to either (A) the value in one field (i.e., column), (B) the concatenation of the values in a plurality of columns, or (C) a computed function of the values in one or more columns. The set of columns used to generate the primary key are represented by a vector herein called the PrimaryKeyColumns vector. There is a "Create_PrimaryKey" function for every database table, represented as follows:

PrimaryKey

= Create_PrimaryKey(BaseTable(RecPtr), PrimaryKeyColumns)

where RecPtr is a pointer to a database table record.

[0032] It is often the case that an application program needs to access a database table in accordance with a set of column values, at least some of which are not included in the primary index. When that is the case, a Create Index procedure can be used to create an efficient alternate access path to the database table by ordering data according to the values in any specified set of columns. That ordering is represented by an "Alternate Index," which is typically implemented as a separate data structure from the associated database table.

[0033] Figure 2B shows the data structure of an alternate index 170. The alternate index 170 includes a file label 172, which includes a compactly stored copy of the catalog information for the index. The alternate index also includes an alternate key B-tree 174 and then a data array 176 organized in rows and columns. The data array has two sets of columns herein called the AltKey columns and the PrimaryKey columns.

[0034] The rows of the data array 176 are called records 178, and each row of the Alternate Index corresponds to one record of the associated database table. Furthermore, each row of the Alternate Index has two fields: one represents the alternate key value for the corresponding database table record, and one represents the Primary Key value for the same database table record.

[0035] The set of columns used to generate the alternate key for a particular database table are represented by a vector herein called the AlternateKeyColumns vector. There is a "Create_AltKey" function for every alternate index of any database table, represented as follows:

AltKey = Create_AltKey(BaseTable(RecPtr), PrimaryKeyColumns)

where RecPtr is a pointer to a database table record.

[0036] For the purposes of the Move Partition, Split Partition and Move Partition Boundary operations, a database index may be viewed in the same way as a database table.

[0037] Figure 2C shows the data structure of an Audit Trail 138-1. The Audit Trail includes a file label 182 and a set of sequentially generated and stored audit entries 184. Each audit entry 184 denotes a database table or index record event such as the addition, deletion or alteration of a specified database table or index record in a specified database table or index.

Database Table Alteration Procedures

[0038] Figure 3 is a conceptual representation of the procedure for modifying a database table or index. The commands for making database table or index alterations are called data definition language (DDL) statements. In the preferred embodiment, the DDL statements used are not changed, except for the use of an "ONLINE" option in the DDL statements to indicate that the DDL operation is to be performed while minimizing the impact of the DDL operation on user transactions. When the ONLINE option is specified, the preferred embodiment of the present invention changes how the SQL catalog manager 130 executes the Move Partition, Split Partition, Move Partition Boundary and Create

Index commands.

[0039] The procedure for modifying a database table's structure begins with a user or operator entering a DDL statement, specifying an alteration in the schema of a specified object 200, which is typically either a database table, an index, or a partition of a database table or index. The specified object is accessed to read its file label, and if the command requires generating a new object 202, the new object is created. If the command involves movement of data between the first object 200 and a second object 202 that already exists, the file label of the second object is also accessed.

[0040] In the first phase of execution of the "online DDL command," a "dirty copy" of data from the first object 200 into the second object 202 is made while user transactions against the first object (and against the second object if it existed prior to the command) are allowed to continue unimpeded by the execution of the DDL command. In particular, a "dirty copy" is made by accessing all the data records of the first object that are the subject of the DDL command, using "browse" access, and generating corresponding records in the second object 202 until the last of the relevant data records in the first object 200 have been accessed.

[0041] Browse access is a form of read access that allows database records to be read through any record locks that may be in place. Thus, browse access will sometimes result in the reading of records which are in the midst of being modified. Also, user transactions against the first object may delete or modify records in the first object while or after they are accessed by the first phase of the DDL command execution. As a result, the records created in the second object 202 may require a certain amount of correction.

[0042] In the second phase of execution of the online DDL command, a partial clean-up of the records in the second object is performed by accessing records in the audit trail 204 associated with the first object 200 and performing corresponding "redo" operations against the second object. During this second phase user transactions against the first object (and against the second object if it existed prior to the command) are allowed to continue unimpeded by the execution of the DDL command.

[0043] More specifically, before the dirty copy in the first phase of the online DDL command is executed, a pointer called AuditPtr is set to the end of the audit trail 204. During the second phase, each audit record beginning with the one referenced by the AuditPtr is inspected. If the audit record is relevant to operations performed on the first object 200, or relevant to the subset of records of the first object that are the subject of the DDL command, then a redo operation is performed against the second object 202 based on the information in the audit record. In the case of a Create Index command, the redo operation uses new audit records that are generated based on the audit entries found in the audit trail because the second object has a very different structure than the first object; in the case of the other three online DDL commands, the redo operation uses the audit record as found in the audit trail except that the audit record is modified prior to the redo operation so as to reference the second object instead of the first object, and the redo operation is performed against the second object.

[0044] The second phase continues until the end of the audit trail is reached.

[0045] In the third phase, a lock is requested against the first object and second object (if it existed prior to the DDL command), and when that lock (or those locks) is (are) granted all user transactions other than browse access transactions are blocked until the third phase is completed. During the third phase, any new audit records in the audit trail 204 associated with the first object 200 are accessed and corresponding "redo" operations are performed against the second object 202. Next, if the first object is a portion of a larger database table or index, a lock against the entire associated database table associated with the first object is obtained while catalog and file labels are updated. In particular, the catalog entries and the file labels associated with the first and second objects are updated to reflect the results of the DDL operation. Furthermore, the catalog entries and file labels of all objects which incorporate schema information made invalid by the DDL operation are also updated. Then the lock on the first and second objects (and the lock, if any, on the associated database table or index) is released, enabling user transactions against the first and second objects to resume.

[0046] Finally, if the DDL command requires deletion of the first object or deletion of a range of records in the first object, that deletion operation is performed in a way that permits concurrent transaction activity.

[0047] It is noted that the implementation of database reconfiguration operations using well established, optimized, database recovery process procedures makes the database reconfiguration operations efficient in terms of computation time and in terms of the computer resources used.

[0048] Appendix 1 lists a pseudocode representation of the Move Partition procedure used by the SQL catalog manager to move a database table partition from one disk to another.

[0049] Appendix 2 lists a pseudocode representation of the Split Partition procedure used by the SQL executor to split a database table partition into two partitions.

[0050] Appendix 3 lists a pseudocode representation of the Move Partition Boundary procedure used by the SQL catalog manager to move rows (i.e., records) of a database table between adjacent partitions (i.e., partitions with adjacent ranges of the table's primary index).

[0051] Appendix 4 lists a pseudocode representation of the Create Index procedure used by the SQL catalog man-

ager to create a new "alternate" index for a database table.

[0052] Appendix 5 lists a pseudocode representation of an alternate embodiment of the Create Index procedure.

[0053] Appendices 6, 7 and 8 lists pseudocode representations of procedures used by the Create Index procedure.

[0054] The pseudocode used in Appendices 1 through 8 is, essentially, a computer language using universal computer language conventions. While the pseudocode employed here has been invented solely for the purposes of this description, it is designed to be easily understandable by any computer programmer skilled in the art.

Move Partition Procedure

[0055] The function of the Move Partition procedure is to move a database table or index, or a specified partition of the table or index, from a first disk location to a second disk location. The procedure is executed in response to a Move Partition command. The second disk location is typically located on a different disk from the first, and often will be located on a different computer node from the first. For the purposes of explaining the Move Partition procedure, it will be assumed that a partition of a database table is being moved, with the understanding that the same steps would be performed for moving an entire unpartitioned database file or for moving an index file or a partition of an index file.

[0056] Referring to Figure 4 and Appendix 1, the steps of the Move Partition procedure are as follows. The database partition to be moved, herein called the Old Partition, is accessed, and a new file for the new database partition, herein called the New Partition, is created on the appropriate disk volume. Furthermore, an audit trail pointer, AuditPtr, is set to point to where the next record in the audit trail for transactions against the database table associated with the Old Partition will be located.

[0057] In the first phase (220) of the Move Partition procedure, the records in the Old Partition are accessed using Browse access (i.e., read through locks), and records are copied from the Old Partition to the New Partition. While records are created in the New Partition, an up-to-date primary index B-tree is maintained for the New Partition. To prevent further repetition, it is noted that for all records created and updated by the online DDL procedures, a corresponding B-tree is updated.

[0058] Transactions by end users against the Old Partition are allowed to continue during this phase of the Move Partition procedure, and those transactions continue to insert, delete and update records in the Old Partition and create new audit trail entries in the audit trail.

[0059] In the second phase (222) of the Move Partition procedure, transactions by end users against the Old Partition are allowed to continue, and those transactions insert, delete and update records in the Old Partition and create new audit trail entries in the audit trail.

[0060] At the beginning of the second phase, an AuditTrail filter is established such that only Audit Trail records that pertain to the Old Partition are received for processing (by an Audit Fixup Process started by the SQL catalog manager). The processing of each audit record that passes the filter is called an "audit fixup".

[0061] Each received audit record that passes the filter is initially processed by modifying the audit record to refer to the NewPartition. Next, the modified audit record is inspected to determine whether the change to the table noted in the audit record is already reflected in the copied records. If so, the audit record is ignored. Otherwise, a "redo" with respect to the modified audit record is performed so as to apply the change noted in the modified audit record to the NewPartition. This process is repeated for each received audit record that passes the filter until the end of the audit trail is reached.

[0062] In the third phase (224) of the Move Partition procedure, a transaction request is made for a file lock on the Old Partition. This is a transaction request that goes in the lock queue. User transactions initiated before the third phase are not affected, but the lock request prevents users from committing transactions initiated after the lock request is made. When the requested lock is granted the remainder of the third phase is performed as a unitary transaction. The first part of the third phase transaction is performing the audit fixup process described above for all audit records referencing the Old Partition (i.e., that pass the filter) that were created after completion of the second phase.

[0063] When processing of the last such audit record in the audit trail is completed, the third phase transaction requests a lock on the entire database table associated with the Old Partition. This prevents all user transactions on the entire database table until the lock is released. When the full table lock is granted, the catalog entry for the Old Partition is deleted and a new catalog entry for the New Partition is created. The file labels and catalog entries for all partitions of the database table are also updated so as to reference the New Partition.

[0064] When all catalog and file label updates have been completed, the lock on the database table is released, enabling user transactions against the database table to resume. Finally, the Old Partition is deleted, the disk space used by the Old Partition file is released, and then the third phase transaction is concluded.

Split Partition Procedure

[0065] The function of the Split Partition procedure is to move a portion of a database table, or a portion of specified

partition of a database table, from a first disk location to a second disk location. The procedure is executed in response to a Split Partition command. The second disk location is typically located on a different disk from the first, and often will be located on a different computer node from the first. For the purposes of explaining the Split Partition procedure, it will be assumed that a partition of a database file is being split into two partitions, with the understanding that the same steps would be performed for splitting a previously unpartitioned database file or for splitting an index file or a partition of an index file.

[0066] Referring to Figure 5 and Appendix 2, the steps of the Split Partition procedure are as follows. The database partition to be split, herein called the Old Partition, is accessed, and a new file for the new database partition, herein called the New Partition, is created on the appropriate disk volume. Furthermore, an audit trail pointer, AuditPtr, is set to point to where the next record in the audit trail for transactions against the database table associated with the Old Partition will be located.

[0067] The primary key ranges associated with the Old Partition and New Partition are as follows. The Old Partition's initial primary key range is called OldRange, and its new, reduced primary key range is called NewRange1. The primary key range for the New Partition is called NewRange2.

[0068] In the first phase (230) of the Split Partition procedure, the records in the Old Partition having a primary key in NewRange2 are accessed using Browse access (i.e., read through locks), and are copied from the Old Partition to the New Partition.

[0069] Transactions by end users against the Old Partition are allowed to continue during this phase of the Split Partition procedure, and those transactions insert, delete and update records in the Old Partition and create new audit trail entries in the audit trail.

[0070] In the second phase (232) of the Split Partition procedure, transactions by end users against the Old Partition are allowed to continue, and those transactions insert, delete and update records in the Old Partition and create new audit trail entries in the audit trail.

[0071] At the beginning of the second phase, an AuditTrail filter is established such that only Audit Trail records that pertain to the Old Partition with a primary key value in NewRange2 are received for processing.

[0072] Each remaining received audit record is initially processed by modifying the audit record to refer to the New Partition. Next, the modified audit record is inspected to determine whether the change to the table noted in the audit record is already reflected in the copied records. If so, the audit record is ignored. Otherwise, a "redo" with respect to the modified audit record is performed so as to apply the change noted in the modified audit record to the New Partition.

This process is repeated for each received audit record that passes the filter until the end of the audit trail is reached.

[0073] In the third phase (234) of the Split Partition procedure, a transaction request is made for a lock on the Old Partition. This is a transaction request that goes in the lock queue. User transactions initiated before the third phase are not affected, but the lock request prevents users from committing transactions initiated after the lock request is made. When the requested lock is granted the remainder of the third phase is performed as a unitary transaction. The first part of the third phase transaction is performing the audit fixup process described above for all audit records referencing the Old Partition that were created after completion of the second phase and for which either the old record image or new record image has a primary key value in NewRange2 (i.e., that pass the filter).

[0074] When processing of the last such audit record in the audit trail is completed, the B-tree for Old Partition is split into two B-trees, one for records having a primary key in NewRange1 and the other for records having a primary key in NewRange2.

[0075] Then the third phase transaction requests a lock on the entire database table associated with the Old Partition. This permits existing user transactions that work on the table to complete, but suspends all new user transactions on the entire database table until the lock is released. When the full table lock is granted, the catalog entry for the Old Partition is deleted and a catalog entry for the New Partition is created. The file labels and catalog entries for all partitions of the database table are also updated so as to reference the New Partition.

[0076] When all catalog and file label updates have been completed, the lock on the database table is released, enabling user transactions against the database table to resume. Finally, an access check is put in place to make application transactions unable to access records in the NewRange2 B-tree of the Old Partition, and then the third phase transaction is concluded.

[0077] All records in the NewRange2 part of the Old Partition are deleted without interfering with concurrent application transactions against the table, and then the access check for the Old Partition is removed.

Move Partition Boundary Procedure

[0078] The function of the Move Partition Boundary procedure is to move records in a specified range of primary key values from a first specified partition of a database table to a second adjacent partition of that database table. The second partition is typically located on a different disk from the first, and often will be located on a different computer node from the first. The procedure is executed in response to a Move Partition Boundary command. For the purposes

of explaining the Move Partition Boundary procedure, it will be assumed that records are being moved between adjacent database file partitions, with the understanding that the same steps would be performed for moving records between adjacent index file partitions.

[0079] Referring to Figure 6 and Appendix 3, the steps of the Move Partition Boundary procedure are as follows. The database partition from which records are to be moved is herein called Partition1 and the database partition to which those records are to be moved is called Partition2. An audit trail pointer, AuditPtr, is set to point to where the next record in the audit trail for transactions against the database table will be located.

[0080] The primary key ranges associated with Partition1 and Partition2 are as follows. Partition1's initial primary key range is called OldRange1, and its new, reduced primary key range is called NewRange1. The initial primary key range for the Partition2 is called OldRange2 and its new expanded key range is called NewRange2. The primary key range of the records to be moved, called MoveRange, is equal to the range of primary key values in NewRange2 that are not in OldRange2.

[0081] In the first phase (240) of the Move Partition Boundary procedure, the records in the Old Partition having a primary key in the MoveRange are accessed using Browse access (i.e., read through locks), and are copied from Partition1 to Partition2.

[0082] Transactions by end users against Partition1 and Partition2 are allowed to continue during this phase of the Move Partition Boundary procedure, and those transactions insert, delete and update records in Partition1 and Partition2 and create new audit trail entries in the audit trail.

[0083] In the second phase (242) of the Move Partition Boundary procedure, transactions by end users against Partition1 and Partition2 are allowed to continue, and those transactions insert, delete and update records in Partition1 and Partition2 and create additional new audit trail entries in the audit trail.

[0084] At the beginning of the second phase, an AuditTrail filter is established such that only Audit Trail records that pertain to the Partition1 are received for processing by the SQL catalog manager. Furthermore, those audit records for which neither the old record image (if any) nor the new record image (if any) have a primary key value in the MoveRange are also ignored.

[0085] Each remaining received audit record is initially processed by modifying the audit record to refer to the Partition2. Next, the modified audit record is inspected to determine whether the change to the table noted in the audit record is already reflected in the copied records. If so, the audit record is ignored. Otherwise, a "redo" with respect to the modified audit record is performed so as to apply the change noted in the modified audit record to Partition2. This process is repeated for each received audit record that passes the filter until the end of the audit trail is reached.

[0086] In the third phase (244) of the Move Partition Boundary procedure, a transaction request is made for a lock on the Old Partition. This is a transaction request that goes in the lock queue. User transactions initiated before the third phase are not affected, but the lock request prevents users from committing transactions initiated after the lock request is made. When the requested lock is granted the remainder of the third phase is performed as a unitary transaction. The first part of the third phase transaction is performing the audit fixup process described above for all audit records referencing Partition1 that were created after completion of the second phase and for which either the old record image or new record image has a primary key value in the MoveRange (i.e., that pass the filter).

[0087] When processing of the last such audit record in the audit trail is completed, the B-tree for Partition1 is split into two B-trees, one for records having a primary key in NewRange1 and the other for records having a primary key in the MoveRange.

[0088] Then the third phase transaction requests a lock on the entire database table associated with Partition1 and Partition2. This permits existing user transactions that work on the table to complete, but suspends all new user transactions on the entire database table until the lock is released. When the full table lock is granted, the catalog entries and file labels for Partition1 and Partition2 are updated to reflect their new primary key ranges. Furthermore, the catalog entries for all partitions of the database table are also updated so as to reference the key ranges of Partition1 and Partition2.

[0089] When all catalog and file label updates have been completed, the lock on the database table is released, enabling user transactions against the database table to be initiated and executed. Finally, an access check is put in place to make application transactions unable to access records in the MoveRange B-tree of the OldPartition, and then the third phase transaction is concluded.

[0090] All records in the MoveRange part of the Old Partition are deleted without interfering with concurrent application transactions against the table, and then the access check for the OldPartition is removed.

Create Index Procedure

[0091] The function of the Create Index procedure is to create an alternate index for a database table, where the alternate index is generally a different function of the database table columns than the primary index for the database table.

[0092] Referring to Figure 6 and Appendices 4, 5, 6 and 7, the steps of the Create Index procedure are as follows. The database object from which records are to be indexed is herein called the Base Table and the alternate index being created is called the NewIndex. The set of columns used to generate the alternate key for the NewIndex are represented by a vector called the AlternateKeyColumns vector. Furthermore, a flag called UniqueIndexFlag is set to True if the NewIndex is to be a unique index, which means that every record in the Base Table must have a unique alternate index value, and is set to False if multiple records in the Base Table are allowed to have the same alternate index value.

[0093] An audit trail pointer, AuditPtr, is set to point to where the next record in the audit trail for transactions against the database table will be located.

[0094] In the first phase (250) of the Create Index procedure, the records in the Base Table are accessed using Browse access (i.e., read through locks), and for each such record a "create new index record" procedure 251 is performed. The "create new index record" procedure 251, which is also used in the later phases of the Create Index procedure, begins by computing primary and alternate index values for a particular record using the index generation functions:

PrimaryKey1 = Create_PrimaryKey(Record, PrimaryKeyColumns)

AltKey1 = Create_AltKey(Record, PrimaryKeyColumns)

[0095] If the UniqueIndexFlag is False, an index record with AltKey1 and PrimaryKey1 as its two fields is stored in the NewIndex.

[0096] If the UniqueIndexFlag is set to True, prior to storing an index record with AltKey1 and PrimaryKey1 as its two fields in the NewIndex, the NewIndex is searched to see if NewIndex already has a record with an alternate index value of AltKey1. If such a record is found in the NewIndex, a potential duplicate index problem exists. To resolve whether a duplicate index problem has been encountered, two tests are performed. First, a repeatable read transaction is performed on the Base Table to see if the Base Table still stores a record with primary and alternate key values of PrimaryKey1 and AltKey1. If such a record is found in the Base Table, the second test is performed.

[0097] The existing NewIndex record with an alternate key value of AltKey1 is read (with a repeatable read operation) to determine the primary key value, PrimaryKey2, stored with it. If this NewIndex record still exists, a second repeatable read transaction is performed on the Base Table to see if the Base Table stores a record with a primary key value of PrimaryKey2 and an alternate key value of AltKey1. If both tests produce positive results, the Create Index procedure is aborted because two database records with identical alternate key values have been encountered. Otherwise, if either test returns a negative result, an index record with AltKey1 and PrimaryKey1 as its two fields is stored in the NewIndex, and the index record with AltKey1 and PrimaryKey2 (if any) is deleted from the NewIndex.

[0098] In the second phase (252) of the Create Index procedure, transactions by end users against the Base Table are allowed to continue, and those transactions insert, delete and update records in the Base Table and create additional new audit trail entries in the audit trail.

[0099] At the beginning of the second phase, an AuditTrail filter is established such that only Audit Trail records that pertain to the Base Table are received for processing by the SQL catalog manager. Each received audit record is processed by an "audit fixup for create index" procedure 253 specially designed for use as part of the Create Index procedure.

[0100] The "audit fixup for create index" procedure 253 processes an audit record as follows. If the audit record denotes a change to a Base Table record that does not alter either the primary key or the alternate key of the record, or the change is already reflected in the NewIndex, the audit record is ignored.

[0101] If the audit record denotes deletion of a Base Table record, an audit record representing deletion of the associated NewIndex record is constructed and a "redo" of the newly created audit record against the NewIndex is performed. As a result, the associated NewIndex record is deleted, if one exists.

[0102] If the audit record denotes addition of a Base Table record, an audit record representing insertion of a corresponding NewIndex record is constructed and a "redo" of the newly created audit record against the NewIndex is performed. If a unique alternate index is being created and a duplicate NewIndex record with the same primary and alternate key values is found during this procedure, the index creation procedure is aborted. Otherwise, the aforementioned steps result in a new index record being stored in the NewIndex for the inserted Base Table record.

[0103] Finally, if the audit record indicates that either the primary index or alternate index of a record have been altered, then a first audit record representing deletion of the NewIndex record associated with the record's old value is constructed, and a second audit record representing insertion of a new NewIndex record corresponding to the record's new value is constructed. Redo's of both constructed audit records are then performed against the NewIndex. If a

unique alternate index is being created and a duplicate NewIndex record with the same primary and alternate key values is found during this procedure, the index creation procedure is aborted. Otherwise, the aforementioned steps result in a new index record being stored in the NewIndex for the altered Base Table record, and deletion of the prior NewIndex record for that Base Table record.

[0104] In the third phase (254) of the Create Index procedure, a transaction request is made for a lock on the Base Table. This is a transaction request that goes in the lock queue. User transactions initiated before the third phase are not affected, but the lock request prevents users from committing transactions initiated after the lock request is made. When the requested lock is granted the remainder of the third phase is performed as a unitary transaction. The first part of the third phase transaction is performing the "audit fixup for create index" process 253 described above for all audit records referencing the Base Table that were created after completion of the second phase (i.e., that pass the filter).

[0105] When processing of the last such audit record in the audit trail is completed, a catalog entry for the NewIndex is created, and file labels for all partitions of the Base Table are updated to reflect the existence of the NewIndex. Then the lock on the Base Table is released, the third phase transaction is concluded, and the Create Index procedure is completed.

[0106] In an alternate embodiment, represented in Appendix 5, the first phase of the Create Index procedure utilizes an indexing procedure that processes the entire Base Table, even if records that potentially have duplicate alternate key values are encountered. All such potential duplicate records are copied to a scratch pad memory area, and then are processed by the "create new index record" procedure 251 to determine if records with duplicate alternate key values in fact exist, and to add additional records to the NewIndex if no duplicate alternate key values are found.

ALTERNATE EMBODIMENTS

[0107] While the present invention has been described with reference to a few specific embodiments, the description is illustrative of the invention and is not to be construed as limiting the invention. Various modifications may occur to those skilled in the art without departing from the scope of the invention as defined by the appended claims.

APPENDIX 1

Pseudocode Representation of Move Partition Procedure

Procedure MovePartition (OldPartition, NewPartition)

/* OldPartition identifies the partition to be moved
 NewPartition identifies the new partition to be created
 BaseTable is the database table associated with the OldPartition */

{

/* Phase 1: Dirty Copy */

Access current database partition (OldPartition) to be moved

Create file for new database partition (NewPartition)

Create file label for the NewPartition

AuditPtr = End of Audit Trail

/* AuditPtr initially points to where the next record in the Audit Trail will
 be located, when it is generated */

While accessing records in the OldPartition, using Browse access (i.e., read
 through locks), copy records from the OldPartition to the NewPartition

Maintain an up-to-date primary index B-tree for each record copied into the
 NewPartition

All transactions by end users continue to insert, delete and update records in
 the OldPartition and to create new audit trail entries

/* Phase 2: Partial Cleanup */

Establish AuditTrail filter: Access only Audit Trail records that pertain to the
 OldPartition

Do Until End of Audit Trail is reached

{

/* Process entry in Audit Trail at AuditPtr */

Modify the Audit Record to refer to the NewPartition

If the change to the table noted in the audit record is already reflected in
 the copied records

{ Ignore Audit Record }

Else

{

Perform a "Redo" with respect to the modified Audit Record so as to
 apply the change noted in the Audit Record to the NewPartition

```

Update B-tree for the NewPartition for each update to the
NewPartition
5      }
Update AuditPtr to Next Audit Record (skip records that do not pertain to
the OldPartition)
      } /* End Do Loop */
10
/* Phase 3: Final Cleanup */
Request and Wait for Lock on the OldPartition
/* This is a transaction request that goes in the lock queue. This does not
15 affect transactions initiated before Phase 3, but prevents users from
committing transactions initiated after the Lock Request is made. */
/* When Lock is granted ... */
20 Do As A Unitary Transaction
{
LastPtr = Last Record in Audit Trail
/* Previously established Audit Trail Filter continued to filter out audit
25 records not applicable to the OldPartition */
Do Until End of AuditPtr reaches LastPtr:
{
30 Modify the Audit Record to refer to the NewPartition
Perform a "Redo" with respect to the modified Audit Record so as to
apply the change noted in the Audit Record to the NewPartition
Update B-tree for the NewPartition for each update to the
35 NewPartition

Update AuditPtr to Next Audit Record (skip records that do not
40 pertain to the OldPartition)
} /* End Do Loop */

Request Lock on the entire Base Table
45 When Lock is granted:
Delete Catalog entry for the OldPartition
Create Catalog entry for the NewPartition
Update file labels and catalog entries for all partitions of BaseTable
50 so as to reference the NewPartition

Release Lock on BaseTable, enabling user transactions against the
55 Base Table to resume execution

```

Delete the OldPartition and release disk space used by the OldPartition

file

} /* End of transaction */

Return

APPENDIX 2

Pseudocode Representation of Split Partition Procedure

Procedure SplitPartition (BaseTable, OldPartition, OldRange, NewRange1,
NewRange2, NewPartition)

/* OldPartition identifies the partition to be split
NewPartition identifies the new partition to be created
OldRange is the range of Primary Key values currently assigned to
the OldPartition
NewRange1 is the new range of Primary Key values assigned to the
OldPartition
NewRange2 is the range of Primary Key values assigned to the
NewPartition to be created
BaseTable is the database table associated with the OldPartition
*/

Access current database partition (OldPartition) to be split

Create file for new database partition (NewPartition)

Create file label for the NewPartition

AuditPtr = End of Audit Trail

/* AuditPtr initially points to where the next record in the Audit Trail will
be located, when it is generated */

/* Phase 1: Dirty Copy */

While accessing records in the OldPartition, using Browse access (i.e., read
through locks), copy records having a Primary Key value in NewRange2
from the OldPartition to the NewPartition

Maintain an up-to-date primary index B-tree for each record copied into the
NewPartition

All transactions by end users continue to insert, delete and update records in
the OldPartition and to create new audit trail entries

/* Phase 2: Partial Cleanup */

Establish AuditTrail filter: Access only Audit Trail records that pertain to the
OldPartition with a Primary Key value in NewRange2

Do Until End of Audit Trail is reached

{


```

/* Process entry in Audit Trail at AuditPtr */
Modify the Audit Record to refer to Partition2
5  If the change to the table noted in the audit record is not already
    reflected in the copied records
    {
10     Perform a "Redo" with respect to the modified Audit Record so as to
        apply the change noted in the Audit Record to Partition2
        Update B-tree for Partition2 for each update to Partition2
    }

15
    Update AuditPtr to Next Audit Record (skip records that do not pertain to
        Partition1 or with a Primary Key value not in NewRange2)
} /* End Do Loop */

20

/* Phase 3: Final Cleanup */
Request Lock on the OldPartition
25 /* This is a transaction request that goes in the lock queue. This does not
    affect transactions initiated before Phase 3, but prevents users from
    committing transactions initiated after the Lock Request is made. */
/* When Lock is granted ... */
30 Do as a Unitary Transaction:
    {
        LastPtr = Last Record in Audit Trail
        /* Previously established Audit Trail Filter continued to filter out audit
35         records not applicable to NewRange2 in the OldPartition */
        Do Until End of AuditPtr reaches LastPtr
            {
40                 Modify the Audit Record to refer to the NewPartition
                 Perform a "Redo" with respect to the modified Audit Record so as to
                     apply the change noted in the Audit Record to the NewPartition
                 Update B-tree for the NewPartition for each update to the
45                 NewPartition
                 Update AuditPtr to Next Audit Record (skip records that do not
                     pertain to NewRange2 in the OldPartition)
50             }
        } /* End Do Loop */

```

55

Split the B-tree for Partition1 into two B-trees, one for records in
 NewRange1 (those remaining in Partition1) and one for the records
 in NewRange2 to be removed from Partition1 (the Prune B-tree).

Request Lock on the entire Base Table

When Lock is granted:

Update Catalog entry for the OldPartition to reflect new index
 boundaries for the OldPartition

Update label for the OldPartition, including putting in place an
 access check that prevents user transaction from accessing
 NewRange2

Create Catalog entry for the NewPartition

Update file labels and catalog entries for all partitions of BaseTable
 to reference the NewPartition

Release Lock on BaseTable, enabling user transactions against the
 Base Table to resume execution

} /* end of unitary transaction for third phase */

/* Phase 4: Background garbage collection */

In background mode:

Delete all records from the OldPartition that are referenced by the Prune
 B-tree (this releases unused space from the OldPartition file)

Remove access check in the OldPartition

Return

APPENDIX 3

5 **Pseudocode Representation of Move Partition Boundary Procedure**

10 **Procedure MovePartitionBoundary (Partition1, OldRange1, NewRange1,
 Partition2, OldRange2, NewRange2)**

15 **/* BaseTable is database table associated with Partition1 and
 Partition2
 Partition1 and Partition2 identify two partitions having adjacent
 assigned Primary Key Ranges
 OldRange1 is the range of Primary Key values currently assigned to
 Partition1
 OldRange2 is the range of Primary Key values currently assigned to
 Partition2
 NewRange1 is the new range of Primary Key values assigned to
25 Partition1
 NewRange2 is the new range of Primary Key values assigned to
 Partition2
30 */**

Access database Partition1 and Partition2

35 **AuditPtr = End of Audit Trail**

**/* AuditPtr initially points to where the next record in the Audit Trail will
 be located, when it is generated */**

40 **Generate MoveRange = range of Primary Key values for records in
 NewRange2 that are not in OldRange2**

/* Phase 1: Dirty Copy */

45 **While accessing records in Partition1 using Browse access, copy records
 having a Primary Key value in MoveRange from Partition1 to Partition2
 Update primary index B-tree for Partition2 for each record copied into
 Partition2. Leave B-tree for Partition1 unchanged.**

50 **All transactions by end users continue (A) to insert, delete and update
 records in Partition1 and Partition2 using old partition range boundaries
 and (B) to create new audit trail entries**

/* Phase 2: Partial Cleanup */

Establish AuditTrail filter: Access only Audit Trail records that pertain to

Partition1 with Primary Key values in MoveRange

Do Until End of Audit Trail is reached

{

/* Process entry in Audit Trail at AuditPtr */

Modify the Audit Record to refer to Partition2

**If the change to the table noted in the audit record is not already
reflected in the copied records**

{

**Perform a "Redo" with respect to the modified Audit Record so as to
apply the change noted in the Audit Record to Partition2**

Update B-tree for Partition2 for each update to Partition2

}

**Update AuditPtr to Next Audit Record (skip records that do not pertain to
Partition1 or with a Primary Key value not in MoveRange)**

} /* End Do Loop */

/* Phase 3: Final Cleanup */

Request Lock on Partition1 and Partition2

**/* This is a transaction request that goes in the lock queue. This does not
affect transactions initiated before Phase 3, but prevents users from
committing transactions initiated after the Lock Request is made. */**

/* When Lock is granted ... */

Do as a Unitary Transaction:

{

LastPtr = Last Record in Audit Trail

**/* Previously established Audit Trail Filter continued to filter out audit
records not applicable to MoveRange in Partition1 */**

Do Until End of AuditPtr reaches LastPtr

{

Modify the Audit Record to refer to Partition2

**Perform a "Redo" with respect to the modified Audit Record so as to
apply the change noted in the Audit Record to Partition2**

Update B-tree for Partition2 for each update to Partition2

**Update AuditPtr to Next Audit Record (skip records that do not
pertain to MoveRange in Partition1)**

} /* End Do Loop */

Split the B-tree for Partition1 into two B-trees, one for records in
 NewRange1 (those remaining in Partition1) and one for the records
 in the MoveRange to be removed from Partition1 (the Prune B-tree).

Request Lock on entire Base Table

When Lock is granted:

Update Catalog entry for Partition1 to reflect new partition key
 boundaries for Partition1

Update label for Partition1, including putting in place an access
 check that prevents user transactions from accessing
 MoveRange

Update Catalog entry for Partition2 to reflect new partition key
 boundaries for Partition2

Update label for Partition2

Update file labels and catalog entries for all partitions of BaseTable
 to reflect the new partition key boundaries for Partition1 and
 Partition2

Release Locks on BaseTable (including Partition1 and Partition2),
 enabling user transactions against the Base Table to resume
 execution

} /* end of unitary transaction for third phase */

/* Phase 4: Background garbage collection */

In background mode:

Delete all records from Partition1 that are referenced by the Prune
 B-tree (this releases unused space from Partition1 file).

Remove access check in Partition1

Return

APPENDIX 4

Pseudocode Representation of Create Index Procedure

CreateIndex Procedure (BaseTable, AltIndexColumns, UniqueIndexFlag,
NewIndex)

/* BaseTable is table for which index is to be created.

AltIndexColumns is a vector listing the table columns that define the
key for the New Index.

UniqueIndexFlag is equal to True if the new index must be a unique
index (i.e., an index for which every base table record has a unique
alternate key value).

NewIndex is the new index to be created

PrimaryIndexColumns is a vector listing the table columns that
define the key for the Primary Index of the Base Table

*/

/* Structure of NewIndex (and every other index, other than the primary
index) is:

File Label - embedded copy of catalog information for index

B-tree - the B-tree for the index

Index Records, where each index record has two fields:

AltKey, PrimaryKey

where the AltKey field stores the value of the AltIndexColumns
for one record of the BaseTable, and the PrimaryKey field
stores the value of the primary key (i.e., the
PrimaryIndexColumns) for that same record

*/

Create File for NewIndex, including root node for the new index's B-tree

Create file label for NewIndex

AuditPtr = End of Audit Trail

/* AuditPtr initially points to where the next record in the Audit Trail will
be located, when it is generated */

RecPtr = First Record of BaseTable

/* Phase 1: Dirty Copy */

Do Until End of BaseTable is Reached

```

5      {
      Access, using Browse access, BaseTable(RecPtr)
      AltKey1 = Create_AltKey( BaseTable(RecPtr), AltIndexColumns)
      PrimaryKey1 = Create_PrimaryKey( BaseTable(RecPtr),
10         PrimaryIndexColumns )

```

/* Check for Duplicate Index Values only if UniqueIndexFlag is True*/

If UniqueIndexFlag

```

15      {
      Search NewIndex for an index record with
      NewIndex.AltKey equal to AltKey1

```

If a match is found

```

20      {
      Call CheckDupRecord (BaseTable, NewIndex,
      BaseTable(RecPtr))
25      If ReturnCode = "fail"
      Abort CreateIndex Procedure
      }

```

Else

```

30      Call CreateIndexRecord (NewIndex, AltKey1, PrimaryKey1)
      }

```

Else /* UniqueIndexFlag is False */

```

35      Call CreateIndexRecord (NewIndex, AltKey1, PrimaryKey1)

```

Increment RecPtr to point to next BaseTable record, if any

```

40      } /* end of Do loop */

```

/* Phase 2: Partial Cleanup */

Establish AuditTrail filter: Access only Audit Trail records that pertain to the Base Table

Do Until End of Audit Trail is reached

```

50      {
      Call AuditRecord_to_NewIndex (AuditPtr, BaseTable, NewIndex,
      IndexPtr) for the current Audit Record
      Update AuditPtr to Next Audit Record (skip records that do not pertain to
55      the Base Table)
      }

```


/* Phase 3: Final Cleanup */

Request Lock on the Base Table

/* This is a transaction request that goes in the lock queue. This does not affect transactions initiated before Phase 3, but prevents users from committing transactions initiated after the Lock Request is made. */

/* When Lock is granted ... */

Do as a Unitary Transaction:

{

LastPtr = Last Record in Audit Trail

/* Previously established Audit Trail Filter continued to filter out audit records not applicable to the Base Table */

Do Until End of AuditPtr = LastPtr

{

Call AuditRecord_to_NewIndex (AuditPtr, BaseTable, NewIndex) for the current Audit Record

Update AuditPtr to Next Audit Record (skip records that do not pertain to the Base Table)

} /* End Do Loop */

Create Catalog entry for NewIndex

Update file labels and catalog entries for all partitions of BaseTable to reference the NewIndex

Release Lock on BaseTable, enabling user transactions against the Base Table to resume execution

} /* end of unitary transaction for third phase */

Return

APPENDIX 5

Pseudocode for Second Preferred Embodiment
of Phase 1 of Create Index Procedure

/* Phase 1: Dirty Copy */

Initialize ScratchPad pointer SP_Ptr

Do Until End of BaseTable is Reached

{

Access, using Browse access, BaseTable(RecPtr)

AltKey1 = Create_AltKey(BaseTable(RecPtr), AltIndexColumns)

**PrimaryKey1 = Create_PrimaryKey(BaseTable(RecPtr),
PrimaryIndexColumns)**

**/* Check for Duplicate Index Values only if UniqueIndex is True*/
If UniqueIndexFlag**

{

**Search NewIndex for an index record with
NewIndex.AltKey equal to AltKey1**

If a match is found

{

Store RecPtr in ScratchPad(SP_Ptr)

Increment SP_Ptr

}

Else

Call CreateIndexRecord (NewIndex, AltKey1, PrimaryKey1)

}

Else /* UniqueIndexFlag is False */

Call CreateIndexRecord (NewIndex, AltKey1, PrimaryKey1)

Increment RecPtr to point to next BaseTable record, if any

} /* end of Do loop */

/* Process Records in ScratchPad */

Initialize ScratchPad pointer SP_Ptr

Do Until End of ScratchPad is reached

{

RecPtr = ScratchPad(SP_Ptr)

Call CheckDupRecord (BaseTable, NewIndex, BaseTable(RecPtr))

```

    If ReturnCode = "fail"
        Abort CreateIndex Procedure
5      Increment SP_Ptr to point to next ScratchPad entry, if any
    }
/* end of Phase 1 */

```

APPENDIX 6

Pseudocode Representation of Create Index Record Procedure

Procedure CreateIndexRecord (NewIndex, NewAltKey, NewPrimaryKey)

```

/*  NewIndex is the new index in which a new record is to be created.
    NewAltKey is the alternate key of a record in the BaseTable for
25      which an index record is to be created.
    NewPrimaryKey is the primary key of a record in the BaseTable for
        which an index record is to be created.
*/

```

Create NewIndexRecord(AltKey,PrimaryKey):

```

35      NewIndexRecord.AltKey = NewAltKey
        NewIndexRecord.PrimaryKey = NewPrimaryKey
    If NewIndexRecord is already present in NewIndex
        Return
40      Else
        {
            Store NewIndexRecord in NewIndex
            Add B-tree entry for the NewIndexRecord
45        }

```

Return

APPENDIX 7

Pseudocode Representation of Check Duplicate Record Procedure

Procedure CheckDupRecord (BaseTable, NewIndex, DupRecord)

```

/* BaseTable is a database table
   NewIndex is an index file being created for the BaseTable
   DupRecord is an index record whose AltKey value is the same as
   the AltKey value of an already existing index record
*/

```

```

/* Generate alternate key value and primary key value from DupRecord */
AltKeyValue2 = Create_AltKey( DupRecord, AltIndexColumns)
PrimaryKeyValue2 = Create_PrimaryKey(DupRecord, PrimaryIndexColumns)

```

As a unitary transaction:

```

{
  Access the base table record, if any, having a primary key value of
  PrimaryKeyValue2, requesting a lock on the requested record
  If Record=BaseTable(PrimaryKeyValue2) is found in base table
  {
    /* Extract the alternate key value for that base record */
    AltKeyValueY = Record(PrimaryKeyValue2).AltKey
    If AltKeyValueY ≠ AltKeyValue2
      TransactionReturn ("DupRecord not found")
  }
  Else
    TransactionReturn ("DupRecord not found")
}

```

If TransactionReturn = "DupRecord not found"

```

/* A Base Table record for DupRecord was not found, indicating that
   an index record for DupRecord is not needed. */
{
  Return (Success)
}

```

```

/* A base table record matching DupRecord's Primary and Alternate keys
   was found in base table.

```

First step: Read the already existing NewIndex record whose alternate key value is DupRecord.AltKey.

Next step: Check to see if a record with the primary key value found in this NewIndex record is still in Base Table

*/

PrimaryKeyValue1 = NewIndex(AltKeyValue2).PrimaryKey

As a unitary transaction:

{

Access the base table record, if any, having a primary key value of

PrimaryKeyValue1, requesting a lock on the requested record

If record is found in base table

{

/* Extract the alternate key value for that base record */

AltKeyValueX = BaseTable(PrimaryKeyValue1).AltKey

If AltKeyValueX ≠ AltKeyValue2

TransactionReturn ("conflicting record not found")

}

Else

TransactionReturn ("conflicting record not found")

}

If TransactionReturn = "conflicting record not found"

{

/* No Duplicate Record Problem because original record has been deleted or altered */

Delete NewIndex(AltKeyValue2)

Call CreateIndexRecord (NewIndex, AltKey2, PrimaryKey2)

Return (Success)

}

Else

{

/* a conflicting base record was found */

Return (Fail, "Duplicate key value found")

}

/* End of CheckDupRecord Procedure */

APPENDIX 8

Pseudocode Representation of "Audit Record to NewIndex" Procedure

Procedure AuditRecord_to_NewIndex (AuditPtr, BaseTable, NewIndex)

/* BaseTable is database table

NewIndex is Index file

AuditPtr is a pointer to an Audit Trail record

*/

/* Ignore irrelevant Audit Record */

If Audit Record denotes a change to a base table record that does not
alter either the PrimaryKey or AltKey of the record

{ Ignore Audit Record }

/* Record deleted from Base Table */

Elseif Audit Record denotes deletion of a base table record

{

Create an audit record representing deletion of the associated
NewIndex record

If this record exists in NewIndex (i.e., matching both the PrimaryKey
and AltKey of the deleted base table record)

Redo the newly created audit record against NewIndex

}

/* New Record Added to Base Table */

Elseif Audit Record denotes addition of a record to the base table

{

Create an audit record representing insertion of the associated
NewIndex record

If UniqueIndexFlag

{

AltKey1 = Create_AltKey(Inserted Record, AltIndexColumns)

Search NewIndex for an index record with

NewIndex.AltKey equal to AltKey1

If a match is found

{

Call CheckDupRecord (BaseTable, NewIndex, Inserted
Record)

```

        If ReturnCode = "fail"
            Abort CreateIndex Procedure
    }
5      Else
        Redo the newly created audit record against NewIndex
    }
10     Else /* UniqueIndexFlag is False */
        Redo the newly created audit record against NewIndex
    }

15 /* PrimaryKey and/or AltKey in Base Table record was altered */
    Elself audit record denotes alteration of a base table record
    {
20        Create a first audit record representing deletion of the NewIndex
            record associated with the base table record's old value
        If this record exists in NewIndex (i.e., matching both the PrimaryKey
            and AltKey of the base table record's old value)
25            Redo the first newly created audit record against NewIndex

        Create a second audit record representing insertion of the NewIndex
30        record associated with the base table record's new (altered)
            value
        If UniqueIndexFlag
        {
35            AltKey1 = Create_AltKey( Altered Record, AltIndexColumns)
            Search NewIndex for an index record with
                NewIndex.AltKey equal to AltKey1
40            If a match is found
                {
                    Call CheckDupRecord (BaseTable, NewIndex, Altered
                        Record)
45                    If ReturnCode = "fail"
                        Abort CreateIndex Procedure
                }
50            Else
                Redo the second newly created audit record against
                    NewIndex
            }
55     Else /* UniqueIndexFlag is False */

```


Redo the second newly created audit record against NewIndex
 } /* end of processing for "table alteration" audit record

Return

Claims

1. A computer system for storing and providing user access to data in stored database objects, comprising:

memory (116) for storing said database objects (120, 200), said database objects comprising database tables, database partitions and database indices wherein said memory resides in a plurality of interconnected computer nodes (102);
 a central processing unit (110);
 a transaction manager (134), coupled to said memory, for managing computational transactions that add, delete and alter data stored in said database objects; said transaction manager (134) including audit trail generation instructions for generating an audit trail (138) and storing said audit trail in said memory, said audit trail including audit records, at least a subset of said audit records each denoting an event selected from the set consisting essentially of addition, deletion and alteration of specified data in a specified one of said database objects;
 a restructuring procedure, executable by said central processing unit, for restructuring a specified one of said database objects;

characterized in that said restructuring procedure includes:

first phase instructions for accessing data in a first of said objects and storing corresponding data in a second of said objects while allowing continued performance of computational transactions against said first object; second phase instructions, for execution after said first phase instructions, for accessing said audit records in said audit trail created by said transaction manager during execution of said first phase instructions and thereafter until execution of said second phase instructions is completed; said second phase instructions updating said data stored in said second object by redoing with respect to said second object each event denoted by said accessed audit trail records; and
 third phase instructions, for execution after said second phase instructions, for obtaining a lock on said first object so as to prevent continued performance of computational transactions against said first object, and for then accessing audit records in said audit trail created by said transaction manager after execution of said second phase instructions; said third phase instructions updating said data in said second object by redoing with respect to said second object each event denoted by said accessed audit trail records.

2. The computer system of claim 1, further **characterized by**:

at least one of said database tables having a plurality of columns, a plurality of records having primary key values based on a first specified subset of said columns, and a primary key index based on said primary key values;
 said restructuring procedure comprising a create index procedure for creating an alternate key index having a unique alternate key value for each record;
 said first object comprising one of said database tables and said second object comprising an alternate key index generated by said create index procedure;
 said first phase instructions including instructions for accessing said records in said one database table and instructions for generating and storing a corresponding index record in said alternate key index, each record in said alternate key index including said primary key value for said corresponding database table record and an alternate key value based on a second specified subset of said columns;
 said second phase and third phase instructions including instructions for deleting an index record in said alternate key index when one of said accessed audit records indicates an event making said alternate key index record invalid, and instructions for adding index records to said alternate key index when said accessed

audit records indicate events that add new database records and events that alter any field of an existing database record corresponding to any column in said first and second specified subsets of columns; and said generating and storing instructions in said first phase instructions, and said adding index records instructions in said second phase and third phase instructions including instructions for determining whether two records in said one database table have identical alternate key index values and for aborting said create index procedure when two such records are found.

3. The computer system of claim 1, further **characterized by**:

at least a subset of said database tables each including a plurality of partitions, each partition storing records having primary key values in a primary key range distinct from all others of said plurality of partitions; each partition stored as a separate data structure in said memory;
 said restructuring procedure comprising a move partition procedure for moving any specified one of said partitions for one of said database tables having a plurality of partitions from a first memory location to a second memory location;
 said first phase instructions comprising instructions for creating a new partition at said second memory location and for copying all records of said specified partition into said new partition, using read only access to said records of said specified partition while performing said copying;
 said second phase instructions update said records in said new partition so as to redo each database table record event denoted by said accessed audit trail records; and
 said third phase instructions comprising instructions, for execution after said second phase instructions, for obtaining a lock on said one database table, then accessing audit records in said audit trail created by said transaction manager after execution of said second phase instructions; said third phase instructions updating said records in said new partition so as to redo each database table record event denoted by said accessed audit trail records.

4. The computer system of claim 3, further **characterized by**:

said one database table having an associated schema denoting information about said one database table's columns; and
 said third phase instructions including instructions for updating said schema information for said one database table to indicate said new partition's memory location, and then releasing said lock on said one database table.

5. A method of storing and providing user access to data in stored database objects, comprising the steps of:

storing said database objects (120, 200) in memory (116), wherein said memory resides in a plurality of interconnected computer nodes (102) and said database objects comprising database tables, partitions and indices;
 while performing computational transactions that add, delete and alter data stored in said database objects, generating an audit trail and storing said audit trail in said memory, said audit trail including audit records, at least a subset of said audit records each denoting an event selected from the set consisting essentially of addition, deletion and alteration of specified data in a specified one of said database objects; and restructuring a specified one of said database objects,

characterized in that said restructuring step includes:

in a first phase, accessing data in a first of said objects and storing corresponding data in a second of said objects while allowing continued performance of computational transactions against said first object;
 in a second phase, accessing said audit records in said audit trail created during said first phase and thereafter until performance of said second phase is completed, and updating said data stored in said second object by redoing with respect to said second object each event denoted by said accessed audit trail records; and
 in a third phase, obtaining a lock on said first object so as to prevent continued performance of computational transactions against said first object, and then accessing audit records in said audit trail generated after said second phase and updating said data in said second object by redoing with respect to said second object each event denoted by said accessed audit trail records.

6. The method of claim 5, further **characterized by**:

at least one of said database tables having a plurality of columns, a plurality of records having primary key values based on a first specified subset of said columns, and a primary key index based on said primary key values;

said restructuring step comprising a step for creating an alternate key index having a unique alternate key value for each record;

said first object comprising one of said database tables and said second object comprising said alternate key index generated by said create index step;

said restructuring step including:

in said first phase, accessing said records in said one database table and generating and storing a corresponding index record in said alternate key index, each record in said alternate key index including said primary key value for said corresponding database table record and an alternate key value based on a second specified subset of said columns;

in said second phase and said third phase, deleting an index record in said alternate key index when one of said accessed audit records indicates an event making said alternate key index record invalid, and adding index records to said alternate key index when said accessed audit records indicate events that add new database records and events that alter any field of an existing database record corresponding to any column in said first and second specified subsets of columns; and

determining whether two records in said one database table have identical alternate key index values and aborting said restructuring step when two such records are found.

7. The method of claim 5, further **characterized by**:

at least a subset of said database tables each including a plurality of partitions, each partition storing records having primary key values in a primary key range distinct from all others of said plurality of partitions; each partition stored as a separate data structure in said memory;

said restructuring step including moving any specified one of said partitions for one of said database tables having a plurality of partitions from a first memory location to a second memory location; said moving step including:

in said first phase, creating a new partition at said second memory location and copying all records of said specified partition into said new partition, using read only access to said records of said specified partition while performing said copying;

in said second phase, accessing audit records in said audit trail created during execution of said first phase and thereafter until performance of said second phase is completed, and updating said records in said new partition so as to redo each database table record event denoted by said accessed audit trail records; and

in said third phase, obtaining a lock on said one database table, then accessing audit records in said audit trail created after completion of said second phase, and updating said records in said new partition so as to redo each database table record event denoted by said accessed audit trail records; said third phase further including releasing said lock on said one database table.

8. The method of claim 7, wherein each database table has an associated schema denoting information about said table's columns,

said third phase of said moving step including updating said schema information for said one database table to indicate said new partition's memory location.

9. A computer-readable memory configured to direct a computer system to store and provide user access to data in stored database objects, and comprising:

a transaction manager (134) for managing computational transactions that add, delete and alter data stored in said database objects; said database objects (120, 200) comprising database tables, partitions and indices; said transaction manager including audit trail generation instructions for generating an audit trail (138), said audit trail including audit records, at least a subset of said audit records each denoting an event selected from the set consisting essentially of addition, deletion and alteration of specified data in a specified one of said database objects; and

a restructuring procedure for restructuring a specified one of said database objects;

characterized in that said restructuring procedure includes:

first phase instructions for accessing data in a first of said objects and storing corresponding data in a second of said objects while allowing continued performance of computational transactions against said first object;
 5 second phase instructions, for execution after said first phase instructions, for accessing said audit records in said audit trail created by said transaction manager during execution of said first phase instructions and thereafter until execution of said second phase instructions is completed; said second phase instructions updating said data stored in said second object by redoing with respect to said second object each event denoted by said accessed audit trail records; and
 10 third phase instructions, for execution after said second phase instructions, for obtaining a lock on said first object so as to prevent continued performance of computational transactions against said first object, and for then accessing audit records in said audit trail created by said transaction manager after execution of said second phase instructions; said third phase instructions updating said data in said second object by redoing with respect to said second object each event denoted by said accessed audit trail records.

15 **10. The computer-readable memory of claim 9, further characterized by:**

at least one of said database tables having a plurality of columns, a plurality of records having primary key values based on a first specified subset of said columns, and a primary key index based on said primary key values;
 20 said restructuring procedure comprising a create index procedure for creating an alternate key index having a unique alternate key value for each record;
 said first object comprising one of said database tables and said second object comprising an alternate key index generated by said create index procedure;
 25 said first phase instructions including instructions for accessing said records in said one database table and instructions for generating and storing a corresponding index record in said alternate key index, each record in said alternate key index including said primary key value for said corresponding database table record and an alternate key value based on a second specified subset of said columns;
 said second phase and third phase instructions including instructions for deleting an index record in said alternate key index when one of said accessed audit records indicates an event making said alternate key index record invalid, and instructions for adding index records to said alternate key index when said accessed audit records indicate events that add new database records and events that alter any field of an existing database record corresponding to any column in said first and second specified subsets of columns; and
 30 said generating and storing instructions in said first phase instructions, and said adding index records instructions in said second phase and third phase instructions including instructions for determining whether two records in said one database table have identical alternate key index values and for aborting said create index procedure when two such records are found.

35 **11. The computer-readable memory of claim 9, further characterized by:**

40 at least a subset of said database tables each including a plurality of partitions, each partition storing records having primary key values in a primary key range distinct from all others of said plurality of partitions; each partition stored as a separate data structure in memory of said computer system;
 said restructuring procedure comprising a move partition procedure for moving any specified one of said partitions for one of said database tables having a plurality of partitions from a first memory location to a second memory location;
 45 said first phase instructions comprising instructions for creating a new partition at said second memory location and for copying all records of said specified partition into said new partition, using read only access to said records of said specified partition while performing said copying;
 50 said second phase instructions update said records in said new partition so as to redo each database table record event denoted by said accessed audit trail records; and
 said third phase instructions comprising instructions, for execution after said second phase instructions, for obtaining a lock on said one database table, then accessing audit records in said audit trail created by said transaction manager after execution of said second phase instructions; said third phase instructions updating said records in said new partition so as to redo each database table record event denoted by said accessed audit trail records.

Patentansprüche

1. Computersystem zum Speichern von und Bereitstellen eines Nutzerzugriffs auf Daten in gespeicherten Datenbankobjekten, umfassend:

einen Speicher (116) zum Speichern der Datenbankobjekte (120, 200), wobei die Datenbankobjekte Datenbanktabellen, Datenbankpartitionen und Datenbankindizes umfassen, wobei der Speicher in einer Mehrzahl von miteinander verbundenen Computerknoten (102) liegt;

eine zentrale Verarbeitungseinheit (110);

einen Transaktionsmanager (134), welcher mit dem Speicher gekoppelt ist, um Computer-Transaktionen zu verwalten, die in den Datenbankobjekten gespeicherte Daten hinzufügen, löschen und verändern; wobei der Transaktionsmanager (134) Prüfpfad-Erzeugungsanweisungen umfasst, um einen Prüfpfad (138) zu erzeugen und um den Prüfpfad in dem Speicher zu speichern, wobei der Prüfpfad Prüfdatensätze umfasst, wobei von wenigstens einem Untersatz der Prüfpfad-Datensätze jeder ein Ereignis bezeichnet, das aus dem Satz ausgewählt ist, der im Wesentlichen aus Hinzufügung, Löschung und Veränderung von spezifizierten Daten in einem spezifizierten der Datenbankobjekte besteht;

eine Restrukturierungsprozedur, welche durch die zentrale Verarbeitungseinheit ausführbar ist, um ein spezifiziertes der Datenbankobjekte zu restrukturieren;

dadurch gekennzeichnet, dass die Restrukturierungsprozedur umfasst:

Erste-Phase-Anweisungen, um auf Daten in einem ersten der Objekte zuzugreifen und um entsprechende Daten in einem zweiten der Objekte zu speichern, während eine fortgesetzte Durchführung von Computer-Transaktionen gegen das erste Objekt gestattet wird;

Zweite-Phase-Anweisungen, zur Ausführung nach den Erste-Phase-Anweisungen, um auf die Prüfdatensätze in dem Prüfpfad zuzugreifen, welche durch den Transaktionsmanager während einer Ausführung der Erste-Phase-Anweisungen und danach, bis eine Ausführung der Zweite-Phase-Anweisungen abgeschlossen ist, erzeugt werden; wobei die Zweite-Phase-Anweisungen die in dem zweiten Objekt gespeicherten Daten aktualisieren, indem sie jedes Ereignis, das durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird, bezüglich des zweiten Objekts wiederholen; sowie

Dritte-Phase-Anweisungen, zur Ausführung nach den Zweite-Phase-Anweisungen, um eine Sperre des ersten Objekts zu erhalten, um eine fortgesetzte Durchführung von Computer-Transaktionen gegen das erste Objekt zu verhindern, sowie um dann auf Prüfdatensätze in dem Prüfpfad zuzugreifen, welche durch den Transaktionsmanager nach einer Ausführung der Zweite-Phase-Anweisungen erzeugt werden; wobei die Dritte-Phase-Anweisungen die Daten in dem zweiten Objekt aktualisieren, indem sie jedes Ereignis, das durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird, bezüglich des zweiten Objekts wiederholen.

2. Computersystem des Anspruchs 1, ferner **dadurch gekennzeichnet:**

dass wenigstens eine der Datenbanktabellen eine Mehrzahl von Spalten aufweist, wobei eine Mehrzahl von Datensätzen Primärschlüsselwerte auf Grundlage eines ersten spezifizierten Untersatzes der Spalten sowie einen Primärschlüsselindex auf Grundlage der Primärschlüsselwerte aufweist;

dass die Restrukturierungsprozedur eine Index-erzeugen-Prozedur umfasst, um einen Alternativschlüsselindex mit einem eindeutigen Alternativschlüsselwert für jeden Datensatz zu erzeugen;

dass das erste Objekt eine der Datenbanktabellen umfasst und das zweite Objekt einen durch die Index-erzeugen-Prozedur erzeugten Alternativschlüsselindex umfasst;

dass die Erste-Phase-Anweisungen Anweisungen umfassen, um auf die Datensätze in der einen Datenbanktabelle zuzugreifen und Anweisungen umfassen, um einen entsprechenden Indexdatensatz in dem Alternativschlüsselindex zu erzeugen und zu speichern, wobei jeder Datensatz in dem Alternativschlüsselindex den Primärschlüsselwert für den entsprechenden Datenbanktabellen-Datensatz und einen Alternativschlüsselwert auf Grundlage eines zweiten spezifizierten Untersatzes der Spalten umfasst;

dass die Zweite-Phase- und Dritte-Phase-Anweisungen Anweisungen umfassen, um einen Indexdatensatz in den Alternativschlüsselindexdaten dann zu löschen, wenn einer der Prüfdatensätze, auf die zugegriffen wird, ein Ereignis anzeigt, das den Alternativschlüsselwert-Datensatz ungültig macht, und Anweisungen umfassen, um Indexdatensätze zu dem Alternativschlüsselindex dann hinzufügen, wenn die Prüfdatensätze, auf die zugegriffen wird, Ereignisse anzeigen, die neue Datenbank-Datensätze hinzufügen, und Ereignisse anzeigen, die irgendein Feld eines bestehenden Datenbank-Datensatzes ändern, entsprechend irgendeiner Spalte in dem ersten und dem zweiten spezifizierten Untersatz von Spalten; und

dass die Erzeugungs- und Speicherungsanweisungen in den Erste-Phase-Anweisungen und die Indexdaten-sätze-hinzufügen-Anweisungen in den Zweite-Phase- und den Dritte-Phase-Anweisungen Anweisungen umfassen, um zu bestimmen, ob zwei Datensätze in der einen Datenbank identische Alternativschlüsselindexwerte aufweisen, und um die Index-erzeugen-Prozedur dann abzubrechen, wenn zwei derartige Datensätze gefunden werden.

3. Computersystem des Anspruchs 1, ferner **gekennzeichnet durch:**

wenigstens einen Untersatz der Datenbanktabellen, wobei jede eine Mehrzahl von Partitionen umfasst und wobei jede Partition Datensätze mit Primärschlüsselwerten in einem Primärschlüsselbereich, welcher von allen anderen aus der Mehrzahl von Partitionen verschieden ist, speichert; wobei jede Partition in dem Speicher als eine gesonderte Datenstruktur gespeichert wird; und ferner **dadurch gekennzeichnet:**

dass die Restrukturierungsprozedur eine Partition-bewegen-Prozedur umfasst, um irgendeine spezifizierte der Partitionen für eine der Datenbanktabellen, welche eine Mehrzahl von Partitionen aufweist, von einem ersten Speicherort zu einem zweiten Speicherort zu bewegen;

dass die Erste-Phase-Anweisungen Anweisungen umfassen, um eine neue Partition bei dem zweiten Speicherort zu erzeugen und um alle Datensätze der spezifizierten Partition in die neue Partition zu kopieren, unter Verwendung von Nur-Lesezugriff auf die Datensätze der spezifizierten Partition, während das Kopieren durchgeführt wird;

dass die Zweite-Phase-Anweisungen die Datensätze in der neuen Partition derart aktualisieren, dass sie jedes Datenbanktabellen-Datensatzereignis wiederholen, welches durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird; und

dass die Dritte-Phase-Anweisungen Anweisungen umfassen, zur Ausführung nach den Zweite-Phase-Anweisungen, um eine Sperre der einen Datenbanktabelle zu erhalten, um dann auf Prüfdatensätze in dem Prüfpfad zuzugreifen, welche durch den Transaktionsmanager nach Ausführung der Zweite-Phase-Anweisungen erzeugt werden; wobei die Dritte-Phase-Anweisungen die Datensätze in der neuen Partition derart aktualisieren, dass sie jedes Datenbanktabellen-Datensatzereignis wiederholen, welches durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird.

4. Computersystem des Anspruchs 3, ferner **dadurch gekennzeichnet:**

dass die eine Datenbanktabelle ein zugeordnetes Schema aufweist, welches Informationen über die Spalten der einen Datenbanktabelle bezeichnet; und

dass die Dritte-Phase-Anweisungen Anweisungen umfassen, um die Schemainformationen für die eine Datenbanktabelle derart zu aktualisieren, dass sie den Speicherort der neuen Partition anzeigen, und um dann die Sperre der einen Datenbanktabelle zu lösen.

5. Verfahren zum Speichern von und Bereitstellen von Nutzerzugriff auf Daten in gespeicherten Datenbankobjekten, welches die folgenden Schritte umfasst:

Speichern der Datenbankobjekte (120, 200) im Speicher (116), wobei der Speicher in einer Mehrzahl von miteinander verbundenen Computerknoten (102) liegt, und wobei die Datenbankobjekte Datenbanktabellen, Partitionen und Indizes umfassen;

Erzeugen eines Prüfpfads und Speichern des Prüfpfads in dem Speicher, während Computer-Transaktionen durchgeführt werden, welche in den Datenbankobjekten gespeicherte Daten hinzufügen, löschen und verändern, wobei der Prüfpfad Prüfdatensätze umfasst, wobei von wenigstens einem Untersatz der Prüfdatensätze jeder ein Ereignis bezeichnet, das ausgewählt ist aus dem Satz, der im Wesentlichen aus Hinzufügung, Löschung und Veränderung von spezifizierten Daten in einem spezifizierten der Datenbankobjekte besteht; und Restrukturieren eines spezifizierten der Datenbankobjekte,

dadurch gekennzeichnet, dass der Restrukturierungsschritt umfasst:

in einer ersten Phase Zugreifen auf Daten in einem ersten der Objekte und Speichern entsprechender Daten in einem zweiten der Objekte, während eine fortgesetzte Durchführung von Computer-Transaktionen gegen das erste Objekt gestattet wird;

in einer zweiten Phase Zugreifen auf die Prüfdatensätze in dem Prüfpfad, welche während der ersten Phase und danach, bis eine Durchführung der zweiten Phase abgeschlossen ist, erzeugt werden, und Aktualisieren

der in dem zweiten Objekt gespeicherten Daten, indem jedes Ereignis, das durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird, bezüglich des zweiten Objekts wiederholt wird; sowie in einer dritten Phase Erhalten einer Sperre des ersten Objekts, um eine fortgesetzte Durchführung von Computer-Transaktionen gegen das erste Objekt zu verhindern, und dann Zugreifen auf Prüfdatensätze in dem Prüfpfad, welche nach der zweiten Phase erzeugt werden, und Aktualisieren der Daten in dem zweiten Objekt, indem jedes Ereignis, das durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird, bezüglich des zweiten Objekts wiederholt wird.

6. Verfahren des Anspruchs 5, ferner **dadurch gekennzeichnet**:

dass wenigstens eine der Datenbanktabellen eine Mehrzahl von Spalten aufweist, wobei eine Mehrzahl von Datensätzen Primärschlüsselwerte auf Grundlage eines ersten spezifizierten Untersatzes der Spalten sowie einen Primärschlüsselindex auf Grundlage der Primärschlüsselwerte aufweist;
dass der Restrukturierungsschritt einen Schritt zum Erzeugen eines Alternativschlüsselindex mit einem eindeutigen Alternativschlüsselwert für jeden Datensatz umfasst;
dass das erste Objekt eine der Datenbanktabellen umfasst und das zweite Objekt den durch den Index-erzeugen-Schritt erzeugten Alternativschlüsselindex umfasst;
wobei der Restrukturierungsschritt umfasst:

in der ersten Phase Zugreifen auf die Datensätze in der einen Datenbanktabelle und Erzeugen und Speichern eines entsprechenden Indexdatensatzes in dem Alternativschlüsselindex, wobei jeder Datensatz in dem Alternativschlüsselindex den Primärschlüsselwert für den entsprechenden Datenbanktabellen-Datensatz und einen Alternativschlüsselwert auf Grundlage eines zweiten spezifizierten Untersatzes der Spalten umfasst;

in der zweiten Phase und der dritten Phase Löschen eines Indexdatensatzes in dem Alternativschlüsselindex dann, wenn einer der Prüfdatensätze, auf die zugegriffen wird, ein Ereignis anzeigt, das den Alternativschlüsselindex-Datensatz ungültig macht, und Hinzufügen von Indexdatensätzen zu dem Alternativschlüsselindex dann, wenn die Prüfdatensätze, auf die zugegriffen wird, Ereignisse anzeigen, die neue Datenbankdatensätze hinzufügen und Ereignisse anzeigen, die irgendein Feld eines bestehenden Datenbankdatensatzes verändern, entsprechend irgendeiner Spalte in dem ersten und dem zweiten spezifizierten Untersatz von Spalten; sowie

Bestimmen, ob zwei Datensätze in der einen Datenbanktabelle identische Alternativschlüsselindexwerte aufweisen, und Abbrechen des Restrukturierungsschritts dann, wenn zwei derartige Datensätze gefunden werden.

7. Verfahren des Anspruchs 5, ferner **gekennzeichnet durch**:

wenigstens einen Untersatz der Datenbanktabellen, wobei jede eine Mehrzahl von Partitionen umfasst und wobei jede Partition Datensätze mit Primärschlüsselwerten in einem Primärschlüsselbereich, der von allen anderen aus der Mehrzahl von Partitionen verschieden ist, speichert; wobei jede Partition als eine gesonderte Datenstruktur in dem Speicher gespeichert wird; und ferner **dadurch gekennzeichnet**:

dass der Restrukturierungsschritt ein Bewegen irgendeiner spezifizierten der Partitionen für eine der Datenbanktabellen, welche eine Mehrzahl von Partitionen aufweist, von einem ersten Speicherort zu einem zweiten Speicherort umfasst; wobei der Bewegungsschritt umfasst:

in der ersten Phase Erzeugen einer neuen Partition bei dem zweiten Speicherort und Kopieren aller Datensätze der spezifizierten Partition in die neue Partition, unter Verwendung von Nur-Lesezugriff auf die Datensätze der spezifizierten Partition, während das Kopieren durchgeführt wird;

in der zweiten Phase Zugreifen auf Prüfdatensätze in dem Prüfpfad, welche während einer Ausführung der ersten Phase und danach, bis eine Durchführung der zweiten Phase abgeschlossen ist, erzeugt werden, und Aktualisieren der Datensätze in der neuen Partition derart, dass jedes Datenbanktabellen-Datensatzereignis wiederholt wird, welches durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird; und

in der dritten Phase Erhalten einer Sperre der einen Datenbanktabelle, dann Zugreifen auf die Prüfdatensätze in dem Prüfpfad, welche nach Vollendung der zweiten Phase erzeugt werden, und Aktualisieren der Datensätze in der neuen Partition derart, dass jedes Datenbanktabellen-Datensatzereignis wiederholt wird, welches durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird; wobei die dritte Phase ferner ein Lösen der Sperre der einen Datenbanktabelle umfasst.

8. Verfahren des Anspruchs 7, wobei jede Datenbanktabelle ein zugeordnetes Schema aufweist, welches Informa-

tionen über die Spalten der Tabelle bezeichnet,

wobei die dritte Phase des Bewegungsschritts ein Aktualisieren der Schemainformationen für die eine Datenbanktabelle umfasst, um den Speicherort der neuen Partition anzuzeigen.

9. Computerlesbarer Speicher, welcher konfiguriert ist, um ein Computersystem derart zu leiten, dass es Daten in gespeicherten Datenbankobjekten speichert und Nutzerzugriff darauf bereitstellt, und welcher umfasst:

einen Transaktionsmanager (134), um Computer-Transaktionen zu verwalten, welche in den Datenbankobjekten gespeicherte Daten hinzufügen, löschen und verändern; wobei die Datenbankobjekte (120, 200) Datenbanktabellen, Partitionen und Indizes umfassen; wobei der Transaktionsmanager Prüfpfad-Erzeugungsanweisungen umfasst, um einen Prüfpfad (138) zu erzeugen, wobei der Prüfpfad Prüfdatensätze umfasst, wobei von wenigstens einem Untersatz der Prüfdatensätze jeder ein Ereignis bezeichnet, das aus dem Satz ausgewählt ist, welcher im Wesentlichen aus Hinzufügung, Löschung und Veränderung von spezifizierten Daten in einem spezifizierten der Datenbankobjekte besteht; sowie eine Restrukturierungsprozedur, um ein spezifiziertes der Datenbankobjekte zu restrukturieren;

dadurch gekennzeichnet, dass die Restrukturierungsprozedur umfasst:

Erste-Phase-Anweisungen, um auf Daten in einem ersten der Objekte zuzugreifen und um entsprechende Daten in einem zweiten der Objekte zu speichern, während eine fortgesetzte Durchführung von Computer-Transaktionen gegen das erste Objekt gestattet wird;

Zweite-Phase-Anweisungen, zur Ausführung nach den Erste-Phase-Anweisungen, um auf die Prüfdatensätze in dem Prüfpfad zuzugreifen, welche durch den Transaktionsmanager während einer Ausführung der Erste-Phase-Anweisungen und danach, bis eine Ausführung der Zweite-Phase-Anweisungen abgeschlossen ist, erzeugt werden; wobei die Zweite-Phase-Anweisungen die in dem zweiten Objekt gespeicherten Daten aktualisieren, indem sie jedes Ereignis, das durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird, bezüglich des zweiten Objekts wiederholen; sowie

Dritte-Phase-Anweisungen, zur Ausführung nach den Zweite-Phase-Anweisungen, um eine Sperre des ersten Objekts zu erhalten, um eine fortgesetzte Durchführung von Computer-Transaktionen gegen das erste Objekt zu verhindern und um dann auf Prüfdatensätze in dem Prüfpfad zuzugreifen, welche durch den Transaktionsmanager nach einer Ausführung der Zweite-Phase-Anweisungen erzeugt werden; wobei die Dritte-Phase-Anweisungen die Daten in dem zweiten Objekt aktualisieren, indem sie jedes Ereignis, das durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird, bezüglich des zweiten Objekts wiederholen.

10. Computerlesbarer Speicher des Anspruchs 9, ferner **dadurch gekennzeichnet:**

dass wenigstens eine der Datenbanktabellen eine Mehrzahl von Spalten aufweist, wobei eine Mehrzahl von Datensätzen Primärschlüsselwerte auf Grundlage eines ersten spezifizierten Untersatzes der Spalten sowie einen Primärschlüsselindex auf Grundlage der Primärschlüsselwerte aufweist;

dass die Restrukturierungsprozedur eine Index-erzeugen-Prozedur umfasst, um einen Alternativschlüsselindex mit einem eindeutigen Alternativschlüsselwert für jeden Datensatz zu erzeugen;

dass das erste Objekt eine der Datenbanktabellen umfasst und das zweite Objekt einen Alternativschlüsselindex umfasst, welcher durch die Index-erzeugen-Prozedur erzeugt wird;

dass die Erste-Phase-Anweisungen Anweisungen umfassen, um auf Datensätze in der einen Datenbanktabelle zuzugreifen und Anweisungen umfassen, um einen entsprechenden Indexdatensatz in dem Alternativschlüsselindex zu erzeugen und zu speichern, wobei jeder Datensatz in dem Alternativschlüsselindex den Primärschlüsselwert für den entsprechenden Datenbanktabellen-Datensatz und einen Alternativschlüsselwert auf Grundlage eines zweiten spezifizierten Untersatzes der Spalten umfasst;

dass die Zweite-Phase- und Dritte-Phase-Anweisungen Anweisungen umfassen, um einen Indexdatensatz in dem Alternativschlüsselindex dann zu löschen, wenn einer der Prüfdatensätze, auf die zugegriffen wird, ein Ereignis anzeigt, das den Alternativschlüsselindex-Datensatz ungültig macht, und Anweisungen umfassen, um Indexdatensätze zu dem Alternativschlüsselindex dann hinzuzufügen, wenn die Prüfdatensätze, auf die zugegriffen wird, Ereignisse anzeigen, die neue Datenbankdatensätze hinzufügen und Ereignisse anzeigen, die irgendein Feld eines bestehenden Datenbank-Datensatzes verändern, entsprechend irgendeiner Spalte in dem ersten und dem zweiten spezifizierten Untersatz von Spalten; und

dass die Erzeugungs- und Speicherungsanweisungen in den Erste-Phase-Anweisungen und die Indexdaten-sätze-hinzufügen-Anweisungen in den Zweite-Phase- und Dritte-Phase-Anweisungen Anweisungen umfassen, um zu bestimmen, ob zwei Datensätze in der einen Datenbanktabelle identische Alternativschlüsselindex-

dexwerte aufweisen und um die Index-erzeugen-Prozedur dann abzuberechnen, wenn zwei derartige Datensätze gefunden werden.

11. Computerlesbarer Speicher des Anspruchs 9, ferner
gekennzeichnet durch:

wenigstens einen Untersatz der Datenbanktabellen, wobei jede eine Mehrzahl von Partitionen umfasst und wobei jede Partition Datensätze mit Primärschlüsselwerten in einem Primärschlüsselwertbereich, der von allen anderen aus der Mehrzahl von Partitionen verschieden ist, speichert; wobei jede Partition als eine gesonderte Datenstruktur im Speicher des Computersystems gespeichert wird; und ferner **dadurch gekennzeichnet**:

dass die Restrukturierungsprozedur eine Partition-bewegen-Prozedur umfasst, um irgendeine spezifizierte der Partitionen für eine der Datenbanktabellen, welche eine Mehrzahl von Partitionen aufweist, von einem ersten Speicherort zu einem zweiten Speicherort zu bewegen;

dass die Erste-Phase-Anweisungen Anweisungen umfassen, um eine neue Partition bei dem zweiten Speicherort zu erzeugen und um alle Datensätze der spezifizierten Partition in die neue Partition zu kopieren, unter Verwendung von Nur-Lesezugriff auf die Datensätze der spezifizierten Partition, während das Kopieren durchgeführt wird;

dass die Zweite-Phase-Anweisungen die Datensätze in der neuen Partition derart aktualisieren, dass sie jedes Datenbanktabellen-Datensatzereignis wiederholen, welches durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird; und

dass die Dritte-Phase-Anweisungen Anweisungen umfassen, zur Ausführung nach den Zweite-Phase-Anweisungen, um eine Sperre der einen Datenbanktabelle zu erhalten, um dann auf Prüfdatensätze in dem Prüfpfad zuzugreifen, welche durch den Transaktionsmanager nach Ausführung der Zweite-Phase-Anweisungen erzeugt werden; wobei die Dritte-Phase-Anweisungen die Datensätze in der neuen Partition derart aktualisieren, dass sie jedes Datenbanktabellen-Datensatzereignis wiederholen, welches durch die Prüfpfad-Datensätze bezeichnet wird, auf die zugegriffen wird.

Revendications

1. Système informatique destiné à mémoriser et à donner, à un utilisateur, accès à des données dans des objets mémorisés de base de données, comprenant :

une mémoire (116) destinée à mémoriser lesdits objets (120, 200) de base de données, lesdits objets de base de données comprenant des tables de base de données, des partitions de base de données et des index de base de données, ladite mémoire résidant dans une pluralité de noeuds informatiques interconnectés (102) ;
une unité centrale de traitement (110) ;

un gestionnaire (134) de transactions, associé à ladite mémoire, destiné à gérer des transactions informatiques qui ajoutent, suppriment et modifient des données mémorisées dans lesdits objets de base de données ; ledit gestionnaire (134) de transactions incluant des instructions de production de journal témoin destiné à engendrer un journal témoin (138) et à mémoriser ledit journal témoin dans ladite mémoire, ledit journal témoin incluant des enregistrements témoins, chacun d'au moins un sous-ensemble desdits enregistrements témoins notant un événement choisi à partir de l'ensemble constitué essentiellement d'addition, de suppression et de modification de données spécifiées dans l'un, spécifié, desdits objets de base de données ;

une procédure de restructuration, exécutable par ladite unité centrale de traitement, destinée à restructurer l'un, spécifié, desdits objets de base de données ;

caractérisé en ce que ladite procédure de restructuration comprend :

des instructions d'une première phase destinées à accéder à des données dans un premier desdits objets et à mémoriser des données correspondantes dans un deuxième desdits objets tout en autorisant la poursuite de l'exécution de transactions informatiques sur ledit premier objet ;

des instructions d'une deuxième phase, pour exécution après lesdites instructions de première phase, destinées à accéder auxdits enregistrements témoins dans ledit journal témoin créé par ledit gestionnaire de transactions pendant l'exécution desdites instructions de première phase et ensuite jusqu'à ce que l'exécution desdites instructions de deuxième phase soit terminée ; lesdites instructions de deuxième phase mettant à jour lesdites données mémorisées dans ledit deuxième objet en refaisant, en ce qui concerne ledit deuxième objet, chaque événement noté par lesdits enregistrements de journal témoin auxquels on a accédé ; et

des instructions d'une troisième phase, pour exécution après lesdites instructions de deuxième phase, destinées à obtenir un verrouillage sur ledit premier objet de façon à prévenir la poursuite de l'exécution de transactions informatiques sur ledit premier objet et à accéder alors aux enregistrements témoins dans ledit journal témoin créé par ledit gestionnaire de transactions après exécution desdites instructions de deuxième phase ;
 5 lesdites instructions de troisième phase mettant à jour lesdites données dans ledit deuxième objet en refaisant, en ce qui concerne ledit deuxième objet, chaque événement noté par lesdits enregistrements de journal témoin auxquels on a accédé.

2. Système informatique selon la revendication 1, caractérisé en outre :

10 en ce qu'au moins l'une desdites tables de base de données comporte plusieurs colonnes, plusieurs enregistrements ayant des valeurs de clé primaire basées sur un premier sous-ensemble spécifié desdites colonnes, et un index de clé primaire basé sur lesdites valeurs de clé primaire ;
 en ce que ladite procédure de restructuration comprend une procédure de création d'index destinée à créer
 15 un index de clé secondaire comportant une unique valeur de clé secondaire pour chaque enregistrement ;
 en ce que ledit premier objet comprend l'une desdites tables de base de données et en ce que ledit deuxième objet comprend un index de clé secondaire engendré par ladite procédure de création d'index ;
 en ce que lesdites instructions de première phase comprennent des instructions pour accéder auxdits enregistrements dans ladite une table de base de données et des instructions pour engendrer et mémoriser un
 20 enregistrement d'index correspondant dans ledit index de clé secondaire, chaque enregistrement dans ledit index de clé secondaire incluant ladite valeur de clé primaire pour ledit enregistrement correspondant de table de base de données et une valeur de clé secondaire basée sur un deuxième sous-ensemble spécifié desdites colonnes ;
 en ce que lesdites instructions de deuxième phase et de troisième phase comprennent des instructions pour
 25 supprimer un enregistrement d'index dans ledit index de clé secondaire lorsque l'un desdits enregistrements témoins auxquels on a accédé indique un événement rendant invalide ledit enregistrement de clé secondaire, et des instructions pour ajouter des enregistrements d'index audit index de clé secondaire lorsque lesdits enregistrements témoins auxquels on a accédé indiquent des événements qui ajoutent de nouveaux enregistrements de base de données et des événements qui modifient une quelconque zone d'un enregistrement de
 30 base de données existant correspondant à une quelconque colonne dans lesdits premier et deuxième sous-ensembles spécifiés de colonnes ; et
 en ce que lesdites instructions de production et de mémorisation dans lesdites instructions de première phase, et lesdites instructions d'ajout d'enregistrements d'index dans lesdites instructions de deuxième phase et de
 35 troisième phase comprennent des instructions pour déterminer si deux enregistrements dans ladite table de base de données ont des valeurs identiques d'index de clé secondaire et pour faire avorter ladite procédure de création d'index lorsque l'on trouve deux de ces enregistrements.

3. Système informatique selon la revendication 1, caractérisé en outre :

40 par au moins un sous-ensemble desdites tables de base de données comprenant chacun une pluralité de partitions, chaque partition mémorisant des enregistrements ayant des valeurs de clé primaire dans une plage de clés primaires distincte de toutes les autres de ladite pluralité de partitions ; chaque partition étant mémorisée comme une structure de données distincte dans ladite mémoire ;
 en ce que ladite procédure de restructuration comprend une procédure de déplacement de partition, d'un
 45 premier emplacement de mémoire à un second emplacement de mémoire, destinée à déplacer l'une quelconque, spécifiée, desdites partitions pour l'une desdites tables de base de données comportant une pluralité de partitions ;
 en ce que lesdites instructions de première phase comprennent des instructions destinées à créer une nouvelle partition au niveau dudit second emplacement de mémoire et à copier tous les enregistrements de ladite
 50 partition spécifiée dans ladite nouvelle partition, en utilisant un accès à lecture seule auxdits enregistrements de ladite partition spécifiée tout en effectuant ladite copie ;
 en ce que lesdites instructions de deuxième phase mettent à jour lesdits enregistrements dans ladite nouvelle partition de façon à refaire chaque événement d'enregistrement de table de base de données noté par lesdits enregistrements de journal témoin auxquels on a accédé ; et
 55 en ce que lesdites instructions de troisième phase comprennent des instructions, pour exécution après lesdites instructions de deuxième phase, pour obtenir un verrouillage de ladite une table de base de données, puis pour accéder auxdits enregistrements dans ledit journal témoin créé par ledit gestionnaire de transactions après exécution desdites instructions de deuxième phase ; lesdites instructions de troisième phase mettant

à jour lesdits enregistrements dans ladite nouvelle partition de façon à refaire chaque événement d'enregistrement de table de base de données noté par lesdits enregistrements de journal témoin auxquels on a accédé.

4. Système informatique selon la revendication 3, caractérisé en outre :

en ce que ladite une table de base de données possède un schéma associé notant de l'information au sujet des colonnes de ladite une table de base de données ; et
en ce que lesdites instructions de la troisième phase comprennent des instructions pour mettre à jour ladite information de schéma pour ladite une table de base de données pour indiquer ledit emplacement de mémoire de ladite nouvelle partition, et pour libérer alors ledit verrouillage de ladite une table de base de données.

5. Procédé destiné à mémoriser et à donner, à un utilisateur, accès à des données dans des objets mémorisés de base de données, comprenant les étapes consistant :

à mémoriser lesdits objets (120, 200) de base de données, dans une mémoire (116), ladite mémoire résidant dans une pluralité de noeuds informatiques interconnectés (102) et lesdits objets de base de données comprenant des tables, des partitions et des index de base de données ;
tout en effectuant des transactions informatiques qui ajoutent, suppriment et modifient des données mémorisées dans lesdits objets de base de données, à engendrer un journal témoin et à mémoriser ledit journal témoin dans ladite mémoire, ledit journal témoin incluant des enregistrements témoins, chacun d'au moins un sous-ensemble desdits enregistrements témoins notant un événement choisi à partir de l'ensemble constitué essentiellement d'addition, de suppression et de modification de données spécifiées dans l'un, spécifié, desdits objets de base de données ;
à restructurer l'un, spécifié, desdits objets de base de données ;

caractérisé en ce que ladite étape de restructuration comprend :

dans une première phase, le fait d'accéder à des données dans un premier desdits objets et de mémoriser des données correspondantes dans un deuxième desdits objets tout en autorisant la poursuite de l'exécution de transactions informatiques sur ledit premier objet ;
dans une deuxième phase, le fait d'accéder auxdits enregistrements témoins dans ledit journal témoin créé pendant ladite première phase et ensuite jusqu'à ce que l'exécution de ladite deuxième phase soit terminée, et le fait de mettre à jour lesdites données mémorisées dans ledit deuxième objet en refaisant, en ce qui concerne ledit deuxième objet, chaque événement noté par lesdits enregistrements de journal témoin auxquels on a accédé ; et
dans une troisième phase, le fait d'obtenir un verrouillage sur ledit premier objet de façon à prévenir la poursuite de l'exécution de transactions informatiques sur ledit premier objet et d'accéder alors aux enregistrements témoins dans ledit journal témoin engendré après ladite deuxième phase, et le fait de mettre à jour lesdites données dans ledit deuxième objet en refaisant, en ce qui concerne ledit deuxième objet, chaque événement noté par lesdits enregistrements de journal témoin auxquels on a accédé.

6. Procédé selon la revendication 5, caractérisé en outre :

en ce qu'au moins l'une desdites tables de base de données comporte plusieurs colonnes, plusieurs enregistrements ayant des valeurs de clé primaire basées sur un premier sous-ensemble spécifié desdites colonnes, et un index de clé primaire basé sur lesdites valeurs de clé primaire ;
en ce que ladite étape de restructuration comprend une étape destinée à créer un index de clé secondaire comportant une unique valeur de clé secondaire pour chaque enregistrement ;
en ce que ledit premier objet comprend l'une desdites tables de base de données et en ce que ledit deuxième objet comprend ledit index de clé secondaire engendré par ladite étape de création d'index ;
en ce que ladite étape de restructuration comprend :

dans ladite première phase, le fait d'accéder auxdits enregistrements dans ladite une table de base de données et d'engendrer et de mémoriser un enregistrement d'index correspondant dans ledit index de clé secondaire, chaque enregistrement dans ledit index de clé secondaire incluant ladite valeur de clé primaire pour ledit enregistrement correspondant de table de base de données et une valeur de clé secondaire basée sur un deuxième sous-ensemble spécifié desdites colonnes ;
dans ladite deuxième phase et dans ladite troisième phase, le fait de supprimer un enregistrement d'index

dans ledit index de clé secondaire lorsque l'un desdits enregistrements témoins auxquels on a accédé indique un événement rendant invalide ledit enregistrement de clé secondaire, et le fait d'ajouter des enregistrements d'index audit index de clé secondaire lorsque lesdits enregistrements témoins auxquels on a accédé indiquent des événements qui ajoutent de nouveaux enregistrements de base de données et des événements qui modifient une quelconque zone d'un enregistrement de base de données existant correspondant à une quelconque colonne dans lesdits premier et deuxième sous-ensembles spécifiés de colonnes ; et le fait de déterminer si deux enregistrements dans ladite table de base de données ont des valeurs identiques d'index de clé secondaire et de faire avorter ladite étape de restructuration lorsque l'on trouve deux de ces enregistrements.

7. Procédé selon la revendication 5, caractérisé en outre :

par au moins un sous-ensemble desdites tables de base de données comprenant chacun une pluralité de partitions, chaque partition mémorisant des enregistrements ayant des valeurs de clé primaire dans une plage de clés primaires distincte de toutes les autres de ladite pluralité de partitions ; chaque partition étant mémorisée comme une structure de données distincte dans ladite mémoire ; en ce que ladite étape de restructuration comprend le fait de déplacer, d'un premier emplacement de mémoire à un second emplacement de mémoire, l'une quelconque, spécifiée, desdites partitions pour l'une desdites tables de base de données comportant une pluralité de partitions ; ladite étape de déplacement incluant :

dans ladite première phase, le fait de créer une nouvelle partition au niveau dudit second emplacement de mémoire et de copier tous les enregistrements de ladite partition spécifiée dans ladite nouvelle partition, en utilisant un accès à lecture seule auxdits enregistrements de ladite partition spécifiée tout en effectuant ladite copie ;

dans ladite deuxième phase, le fait d'accéder auxdits enregistrements témoins dans ledit journal témoin créé pendant l'exécution de ladite première phase et ensuite jusqu'à ce que l'exécution de ladite deuxième phase soit terminée, et le fait de mettre à jour lesdits enregistrements dans ladite nouvelle partition de façon à refaire chaque événement d'enregistrement de table de base de données noté par lesdits enregistrements de journal témoin auxquels on a accédé ; et

dans ladite troisième phase, le fait d'obtenir un verrouillage de ladite une table de base de données, puis d'accéder auxdits enregistrements dans ledit journal témoin créé après l'achèvement de ladite deuxième phase ; et le fait de mettre à jour lesdits enregistrements dans ladite nouvelle partition de façon à refaire chaque événement d'enregistrement de table de base de données noté par lesdits enregistrements de journal témoin auxquels on a accédé ; ladite troisième phase comprenant en outre la libération dudit verrouillage de ladite une table de base de données.

8. Procédé selon la revendication 7, dans lequel chaque table de base de données possède un schéma associé notant de l'information au sujet des colonnes de ladite table ; et

ladite troisième phase de ladite étape de déplacement comprenant le fait de mettre à jour ladite information de schéma pour ladite une table de base de données pour indiquer ledit emplacement de mémoire de ladite nouvelle partition.

9. Mémoire lisible en ordinateur configurée pour diriger un système informatique pour mémoriser et donner, à un utilisateur, accès à des données dans des objets mémorisés de base de données, et comprenant :

un gestionnaire (134) de transactions destiné à gérer des transactions informatiques qui ajoutent, suppriment et modifient des données mémorisées dans lesdits objets de base données ; lesdits objets (120, 200) de base de données comprenant des tables, des partitions et des index de base de données ; ledit gestionnaire de transactions incluant des instructions de production de journal témoin destinées à engendrer un journal témoin (138), ledit journal témoin incluant des enregistrements témoins, chacun d'au moins un sous-ensemble desdits enregistrements témoins notant un événement choisi à partir de l'ensemble constitué essentiellement d'addition, de suppression et de modification de données spécifiées dans l'un, spécifié, desdits objets de base de données ;

une procédure de restructuration, destinée à restructurer l'un, spécifié, desdits objets de base de données ;

caractérisé en ce que ladite procédure de restructuration comprend :

des instructions d'une première phase destinées à accéder à des données dans un premier desdits objets et à mémoriser des données correspondantes dans un deuxième desdits objets tout en autorisant la poursuite de l'exécution de transactions informatiques sur ledit premier objet :

des instructions d'une deuxième phase, pour exécution après lesdites instructions de première phase, destinées à accéder auxdits enregistrements témoins dans ledit journal témoin créé par ledit gestionnaire de transactions pendant l'exécution desdites instructions de première phase et ensuite jusqu'à ce que l'exécution desdites instructions de deuxième phase soit terminée ; lesdites instructions de deuxième phase mettant à jour lesdites données mémorisées dans ledit deuxième objet en refaisant, en ce qui concerne ledit deuxième objet, chaque événement noté par lesdits enregistrements de journal témoin auxquels on a accédé ; et des instructions d'une troisième phase, pour exécution après lesdites instructions de deuxième phase, destinées à obtenir un verrouillage sur ledit premier objet de façon à prévenir la poursuite de l'exécution de transactions informatiques sur ledit premier objet et à accéder alors aux enregistrements témoins dans ledit journal témoin créé par ledit gestionnaire de transactions après exécution desdites instructions de deuxième phase ; lesdites instructions de troisième phase mettant à jour lesdites données dans ledit deuxième objet en refaisant, en ce qui concerne ledit deuxième objet, chaque événement noté par lesdits enregistrements de journal témoin auxquels on a accédé.

10. Mémoire lisible en ordinateur selon la revendication 9, caractérisée en outre :

en ce qu'au moins l'une desdites tables de base de données comporte plusieurs colonnes, plusieurs enregistrements ayant des valeurs de clé primaire basées sur un premier sous-ensemble spécifié desdites colonnes, et un index de clé primaire basé sur lesdites valeurs de clé primaire ; en ce que ladite procédure de restructuration comprend une procédure de création d'index destinée à créer un index de clé secondaire comportant une unique valeur de clé secondaire pour chaque enregistrement ; en ce que ledit premier objet comprend l'une desdites tables de base de données et en ce que ledit deuxième objet comprend un index de clé secondaire engendré par ladite procédure de création d'index ; en ce que lesdites instructions de première phase comprennent des instructions pour accéder auxdits enregistrements dans ladite une table de base de données et des instructions pour engendrer et mémoriser un enregistrement d'index correspondant dans ledit index de clé secondaire, chaque enregistrement dans ledit index de clé secondaire incluant ladite valeur de clé primaire pour ledit enregistrement correspondant de table de base de données et une valeur de clé secondaire basée sur un deuxième sous-ensemble spécifié desdites colonnes ; en ce que lesdites instructions de deuxième phase et de troisième phase comprennent des instructions pour supprimer un enregistrement d'index dans ledit index de clé secondaire lorsque l'un desdits enregistrements témoins auxquels on a accédé indique un événement rendant invalide ledit enregistrement de clé secondaire, et des instructions pour ajouter des enregistrements d'index audit index de clé secondaire lorsque lesdits enregistrements témoins auxquels on a accédé indiquent des événements qui ajoutent de nouveaux enregistrements de base de données et des événements qui modifient une quelconque zone d'un enregistrement de base de données existant correspondant à une quelconque colonne dans lesdits premier et deuxième sous-ensembles spécifiés de colonnes ; et en ce que lesdites instructions de production et de mémorisation dans lesdites instructions de première phase, et lesdites instructions d'ajout d'enregistrements d'index dans lesdites instructions de deuxième phase et de troisième phase comprennent des instructions pour déterminer si deux enregistrements dans ladite table de base de données ont des valeurs identiques d'index de clé secondaire et pour faire avorter ladite procédure de création d'index lorsque l'on trouve deux de ces enregistrements.

11. Mémoire lisible en ordinateur selon la revendication 9, caractérisée en outre :

par au moins un sous-ensemble desdites tables de base de données comprenant chacun une pluralité de partitions, chaque partition mémorisant des enregistrements ayant des valeurs de clé primaire dans une plage de clés primaires distincte de toutes les autres de ladite pluralité de partitions ; chaque partition étant mémorisée comme une structure de données distincte dans la mémoire dudit système informatique ; en ce que ladite procédure de restructuration comprend une procédure de déplacement de partition, d'un premier emplacement de mémoire à un second emplacement de mémoire, destinée à déplacer l'une quelconque, spécifiée, desdites partitions pour l'une desdites tables de base de données comportant une pluralité de partitions ; en ce que lesdites instructions de première phase comprennent des instructions destinées à créer une nouvelle partition au niveau dudit second emplacement de mémoire et à copier tous les enregistrements de ladite

partition spécifiée dans ladite nouvelle partition, en utilisant un accès à lecture seule auxdits enregistrements de ladite partition spécifiée tout en effectuant ladite copie ;

en ce que lesdites instructions de deuxième phase mettent à jour lesdits enregistrements dans ladite nouvelle partition de façon à refaire chaque événement d'enregistrement de table de base de données noté par lesdits enregistrements de journal témoin auxquels on a accédé ; et

en ce que lesdites instructions de troisième phase comprennent des instructions, pour exécution après lesdites instructions de deuxième phase, pour obtenir un verrouillage de ladite une table de base de données, puis pour accéder auxdits enregistrements dans ledit journal témoin créé par ledit gestionnaire de transactions après exécution desdites instructions de deuxième phase ; lesdites instructions de troisième phase mettant à jour lesdits enregistrements dans ladite nouvelle partition de façon à refaire chaque événement d'enregistrement de table de base de données noté par lesdits enregistrements de journal témoin auxquels on a accédé.




FIGURE 1




FIGURE 2A

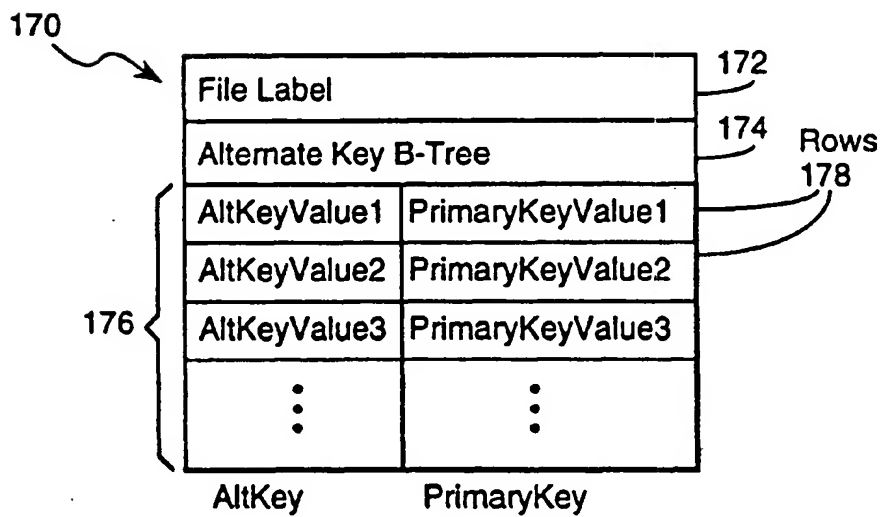


FIGURE 2B

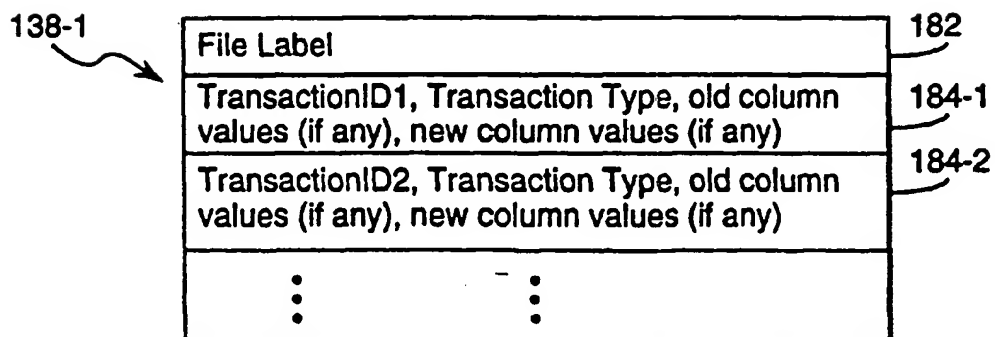



FIGURE 2C

**FIGURE 3**

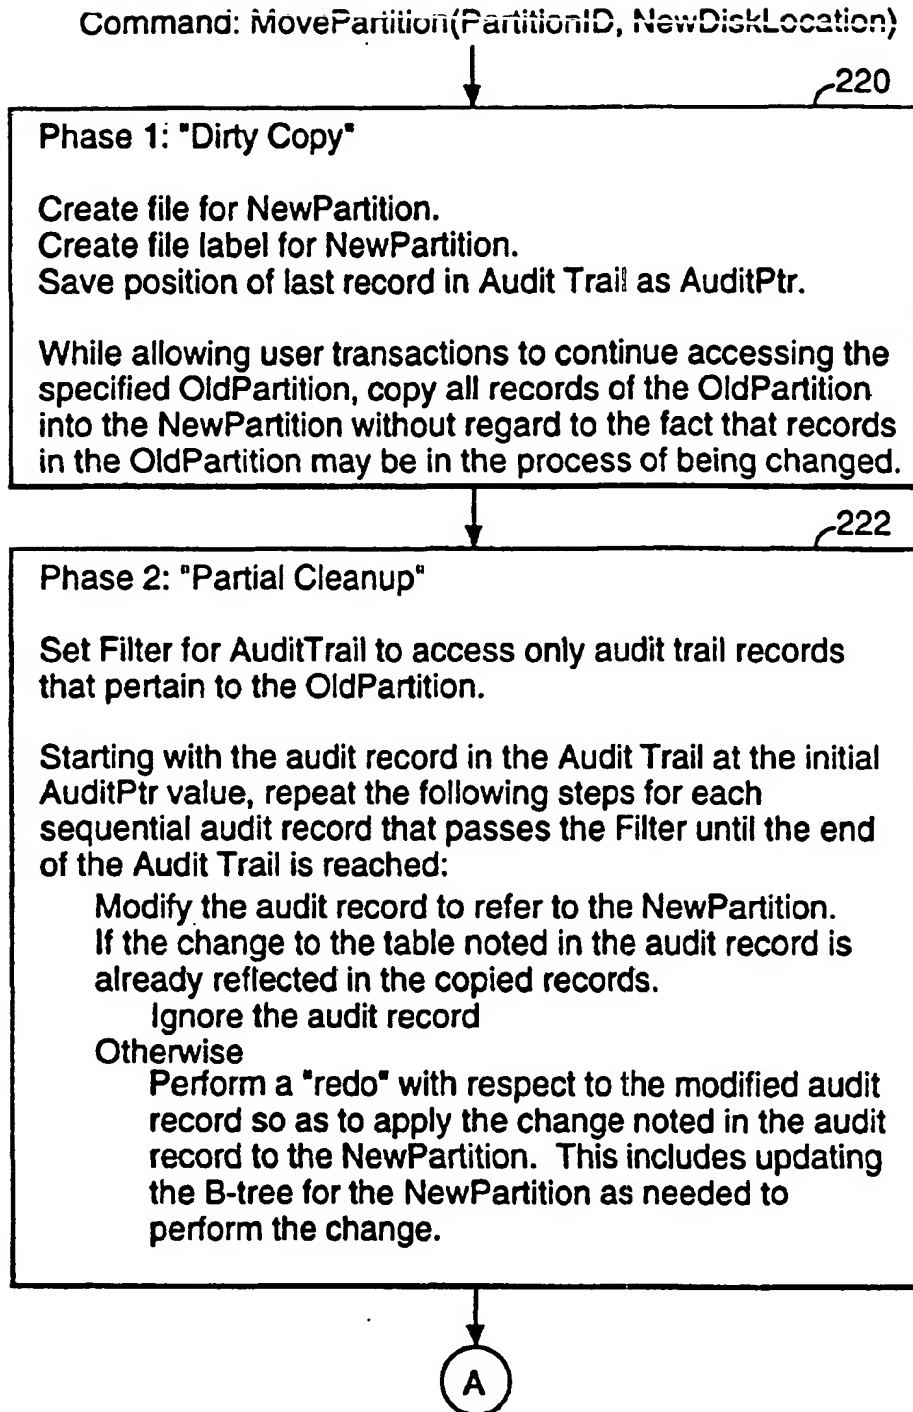
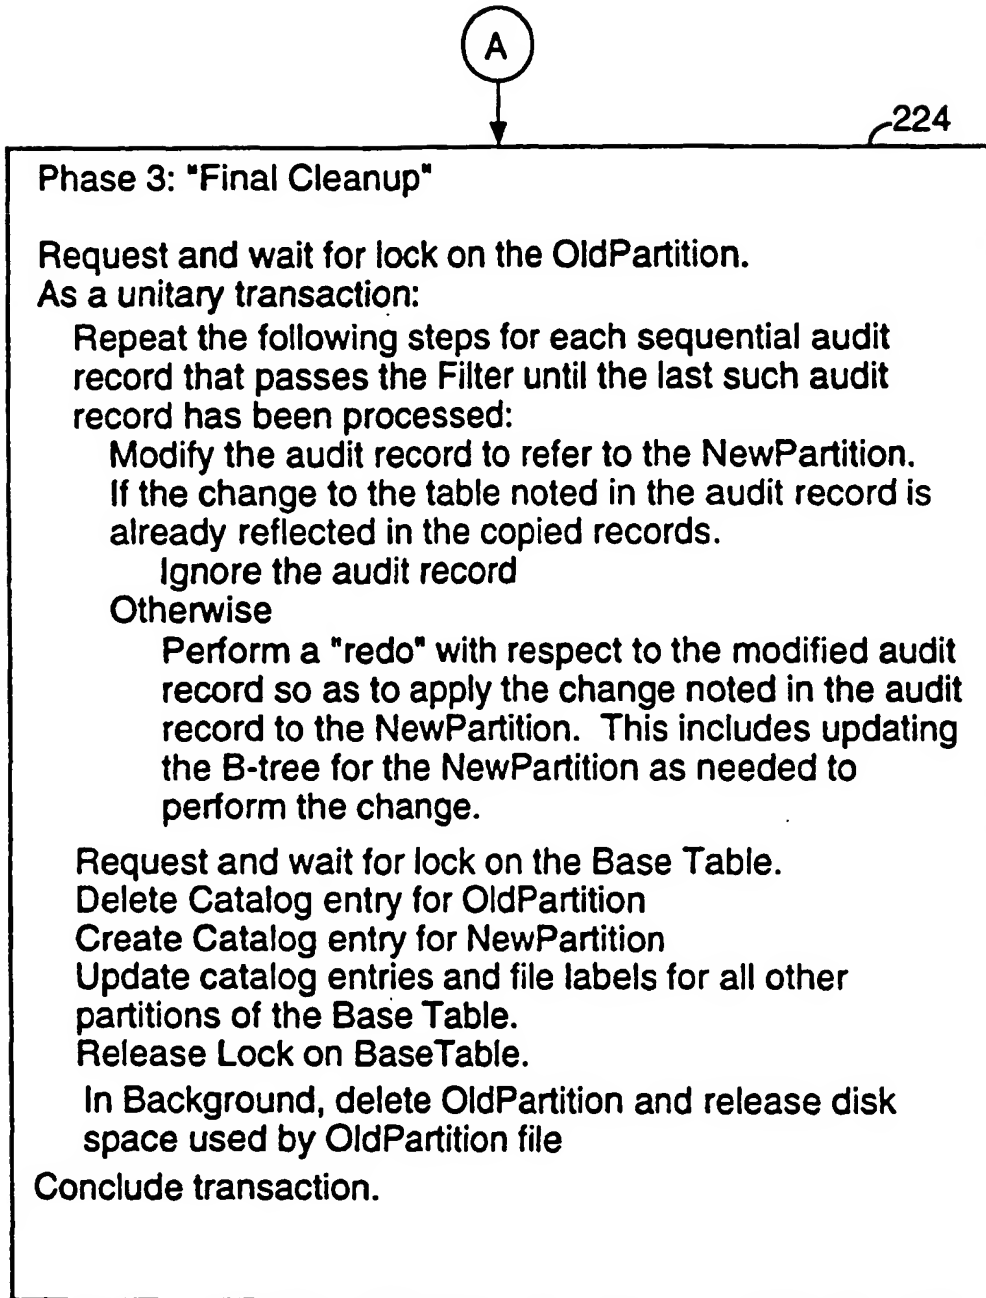


FIGURE 4A

FIGURE 4A

FIGURE 4B

FIGURE 4

**FIGURE 4B**

Command: SplitPartition(OldPartition, OldRange,
NewRange1, NewRange2, NewPartition)

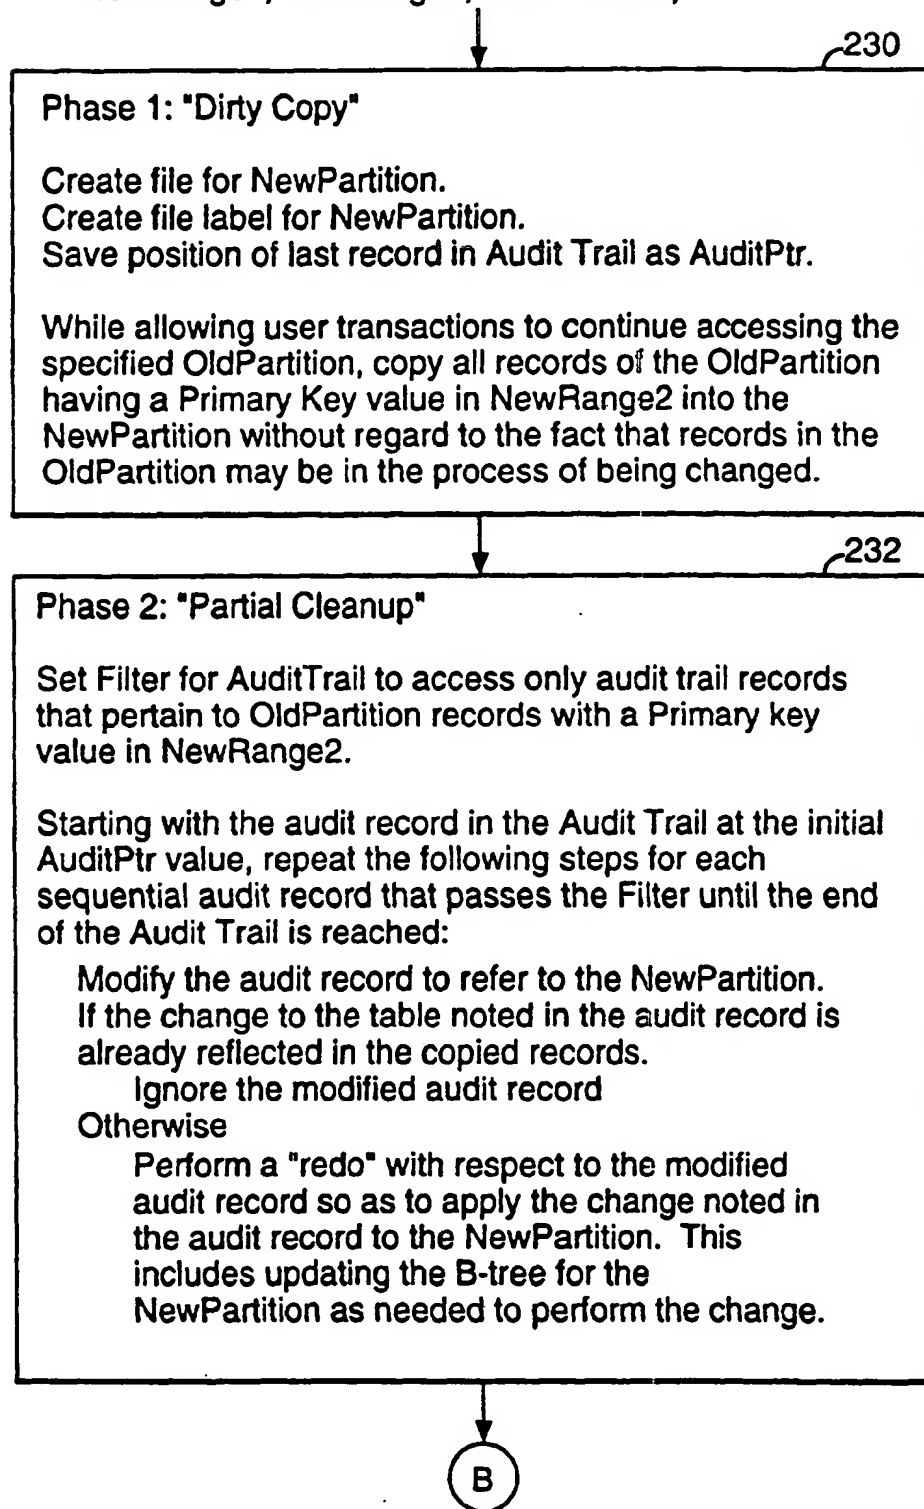
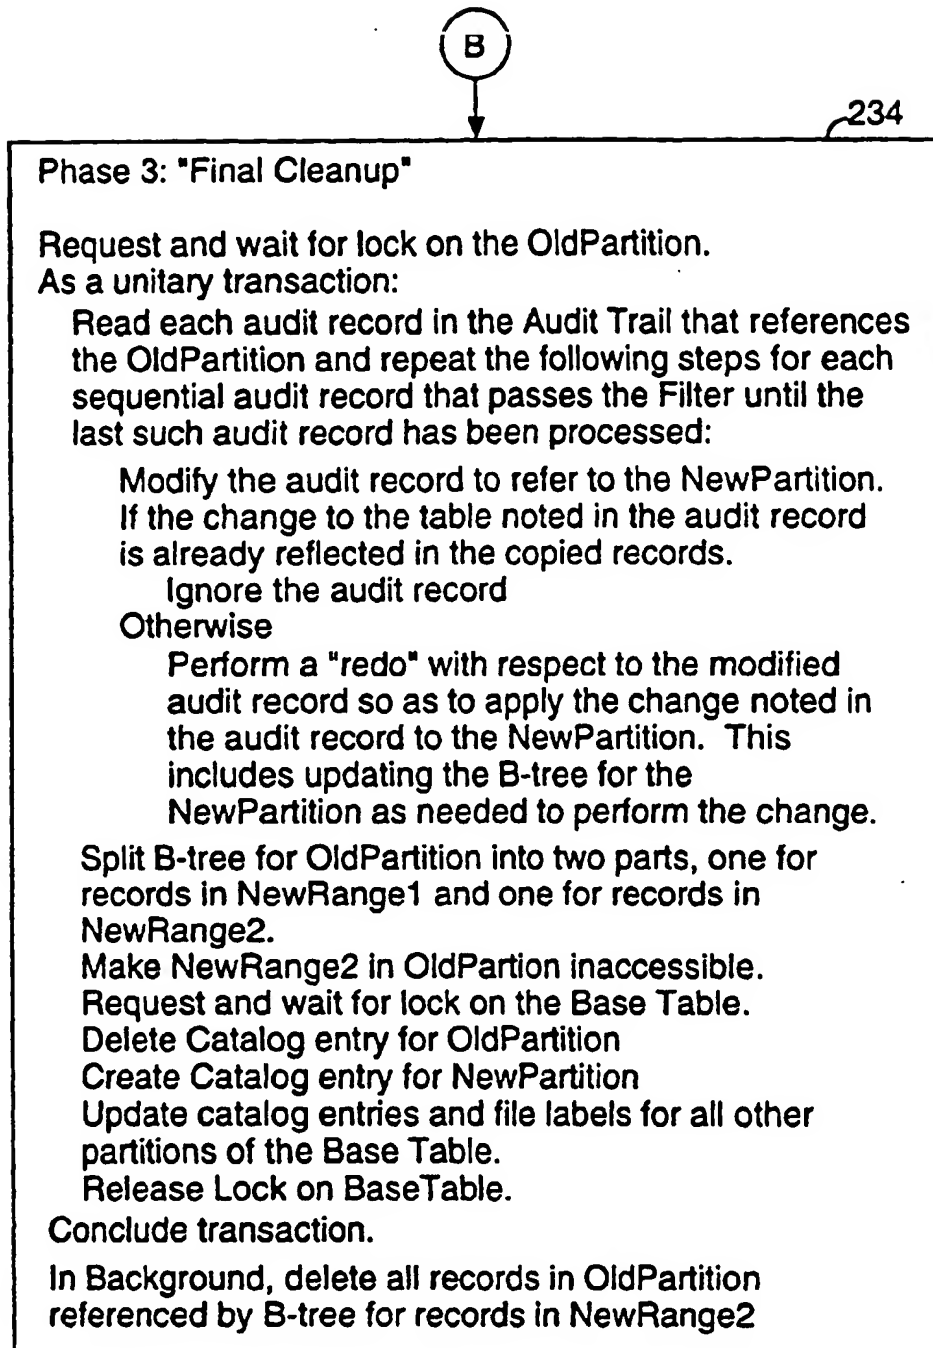
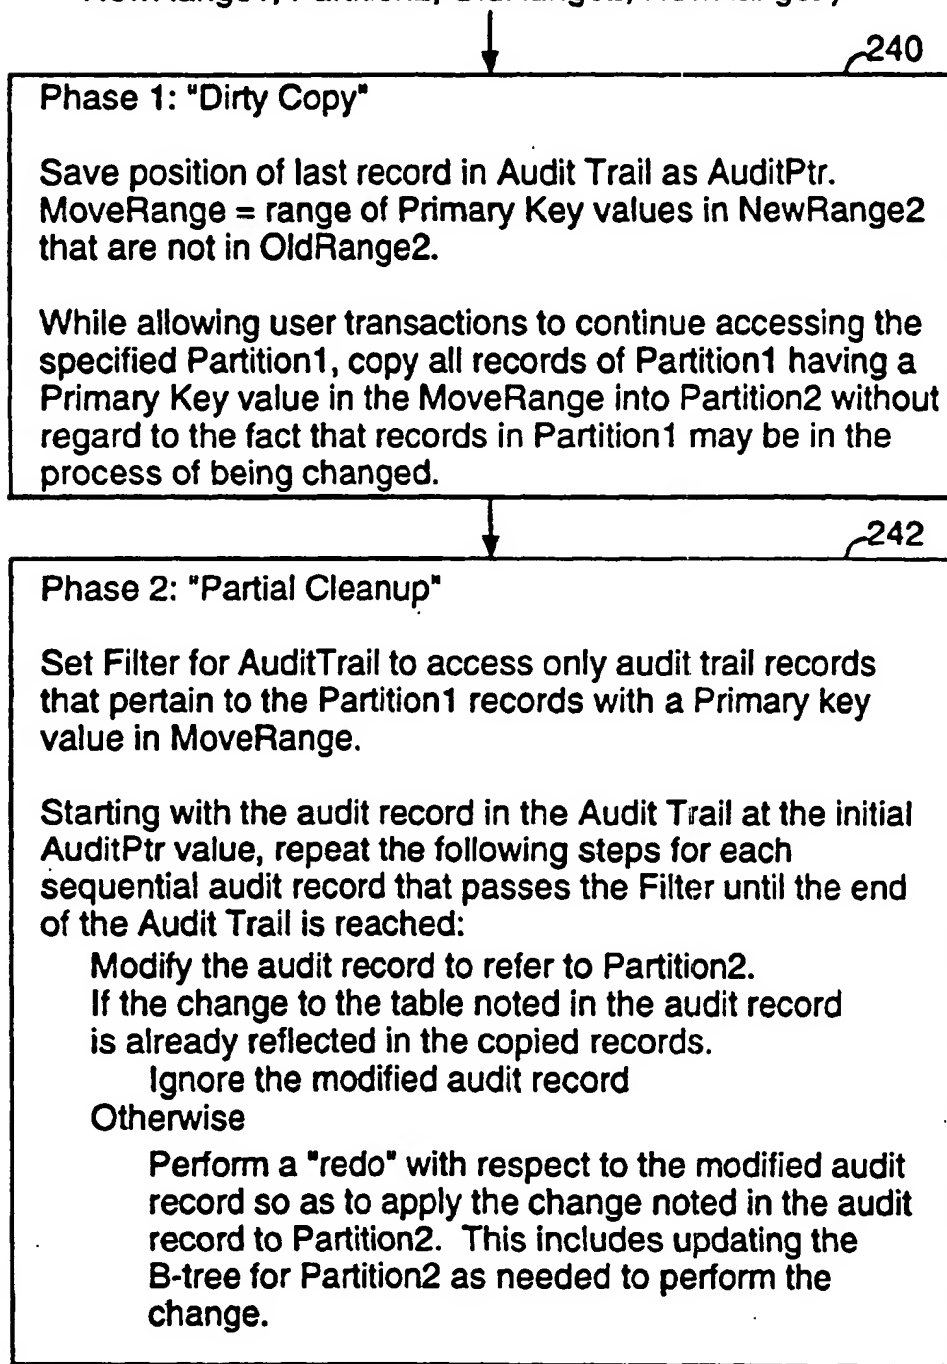


FIGURE 5A

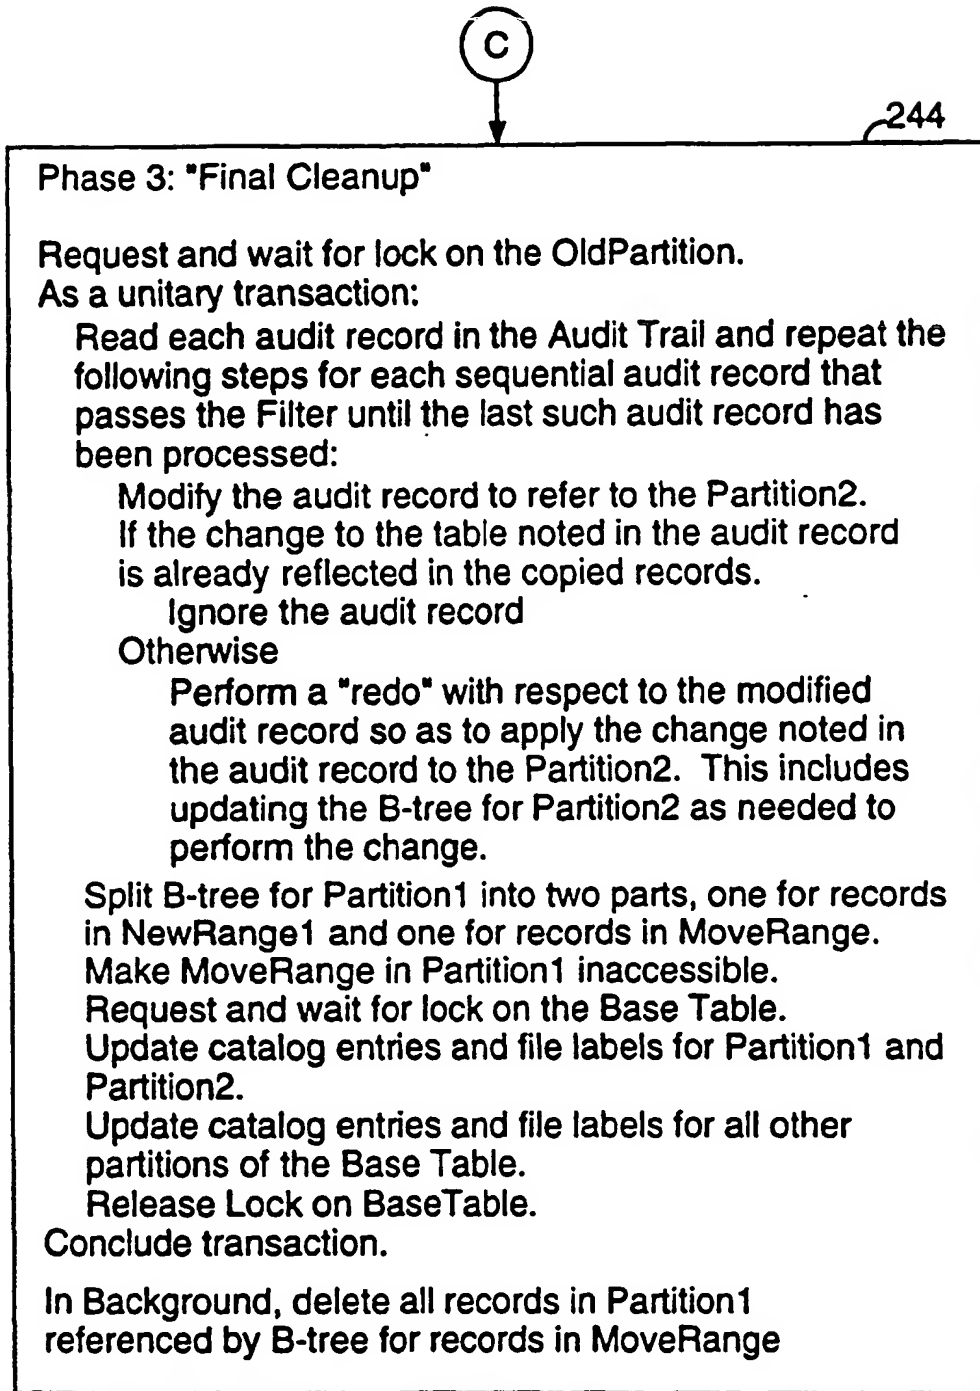
**FIGURE 5B****FIGURE 5A****FIGURE 5B****FIGURE 5 -**

Command: MovePartitionBoundary(Partition1, OldRange1, NewRange1, Partition2, OldRange2, NewRange2)



C

FIGURE 6A

**FIGURE 6B****FIGURE 6A****FIGURE 6B****FIGURE 6**

Command: CreateIndex(BaseTable, AltKeyColumns,
UniqueIndexFlag, NewIndex)

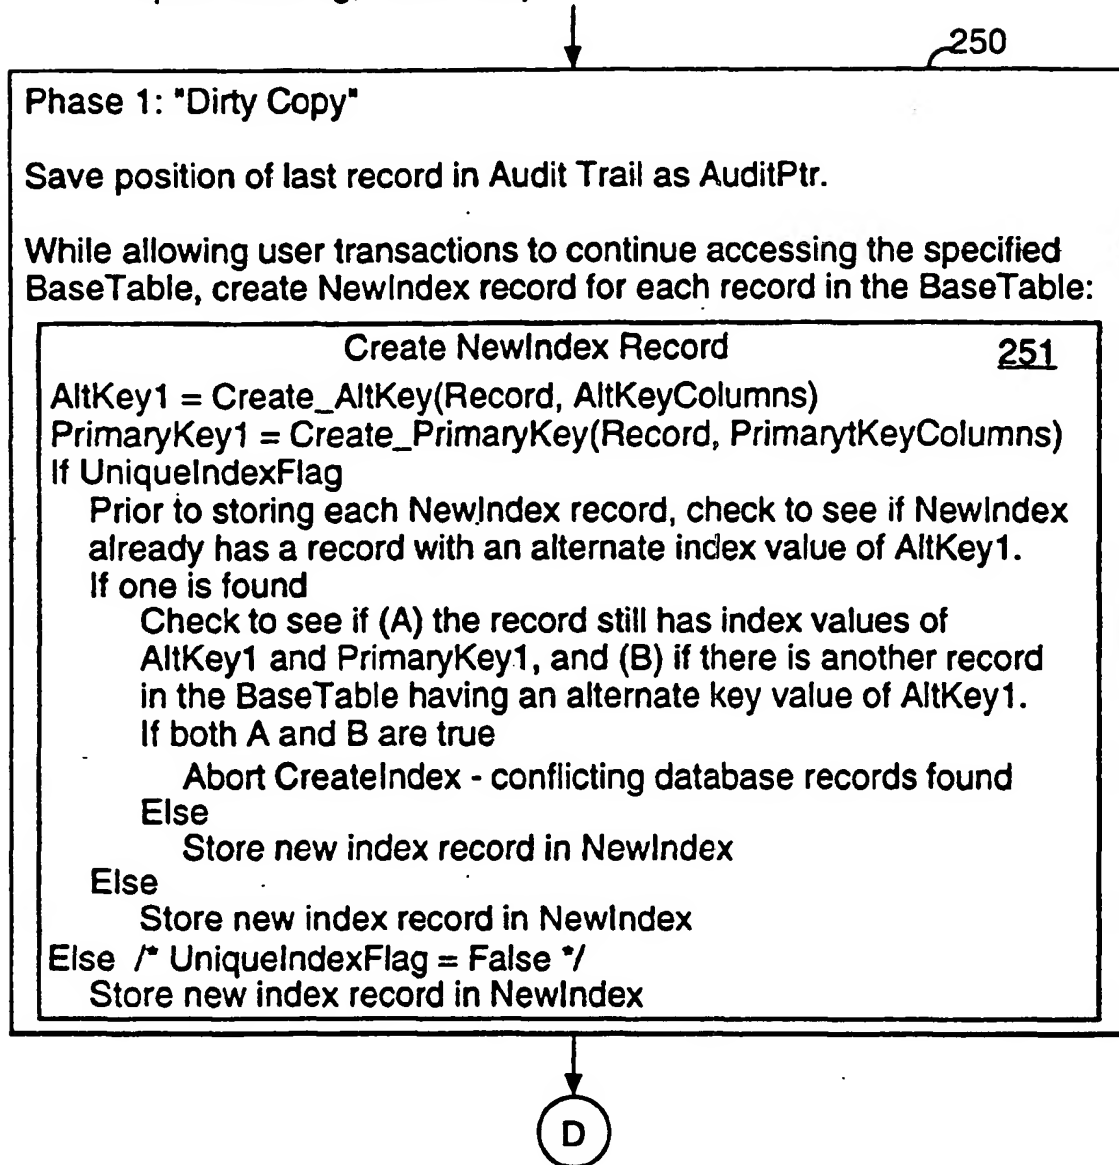


FIGURE 7A

FIGURE 7A
FIGURE 7B
FIGURE 7C

FIGURE 7

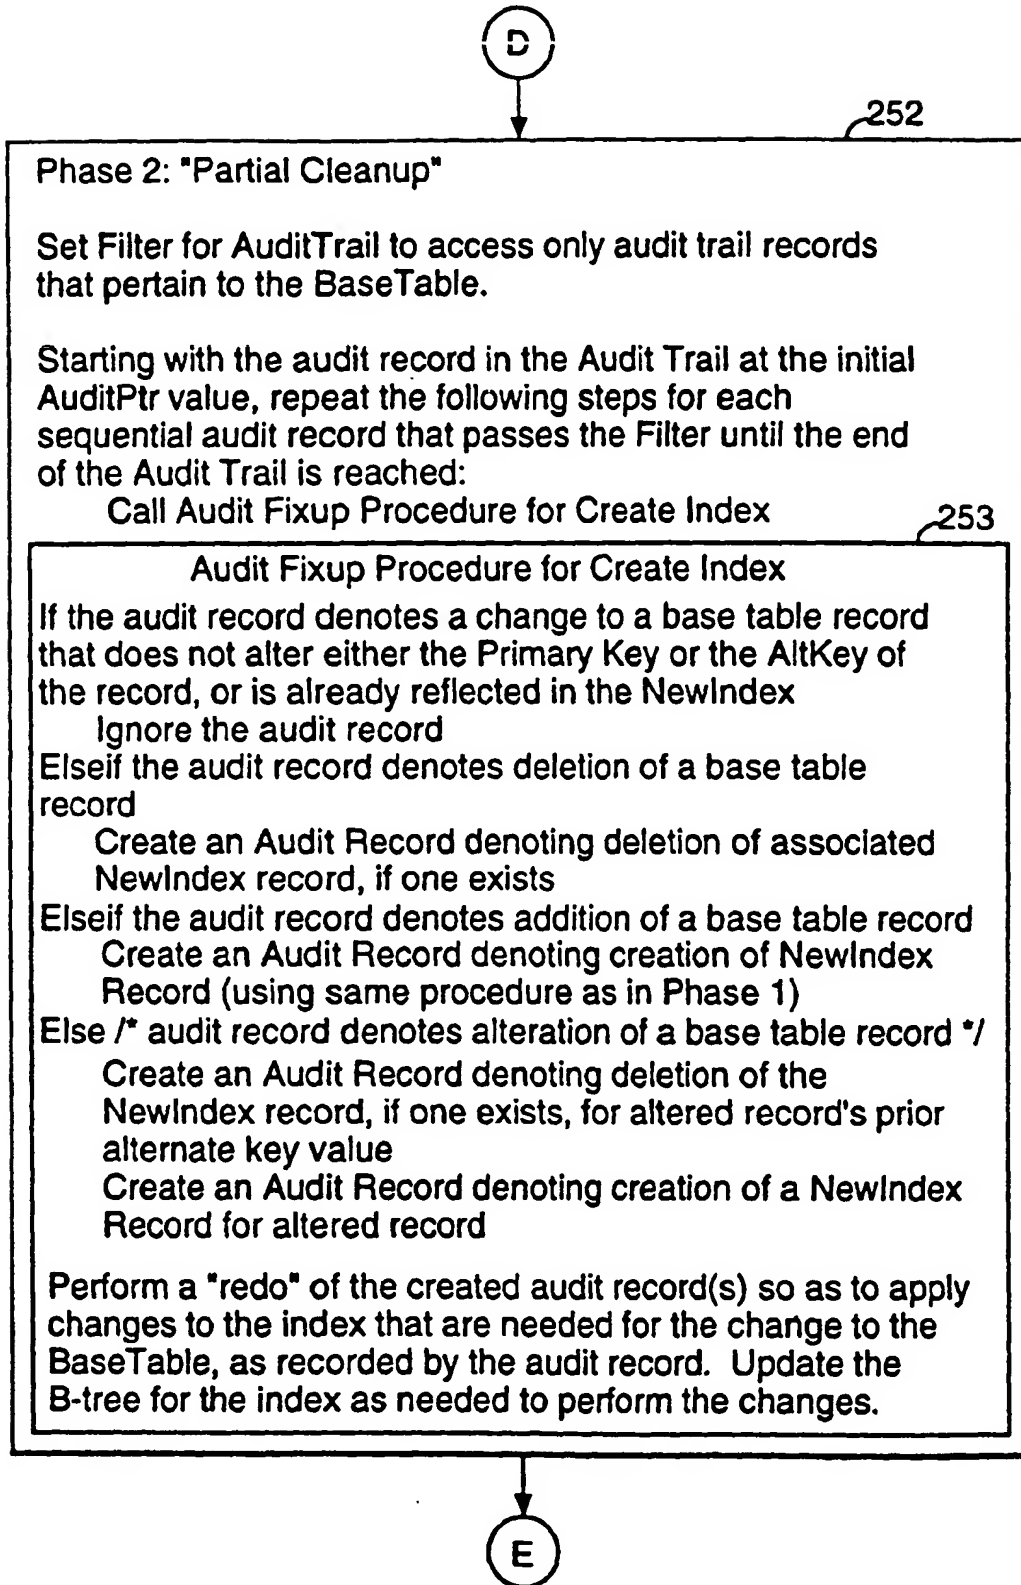


FIGURE 7B

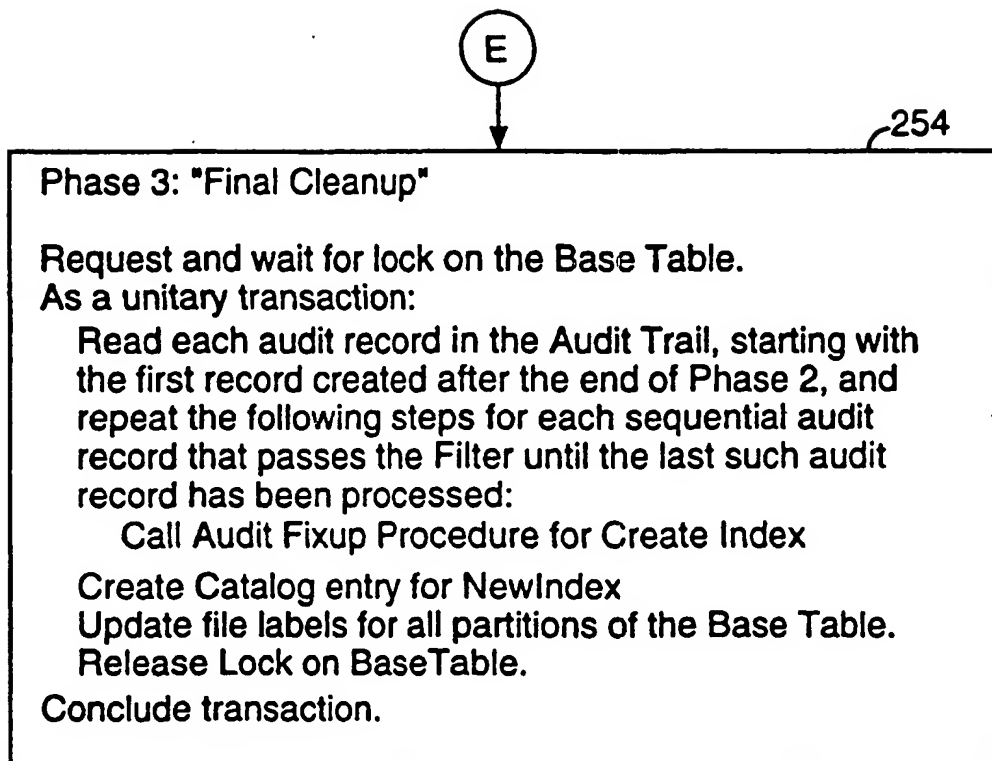


FIGURE 7C